

HP 24-Port 4x Fabric Copper Switch User Guide



November 2004 (First Edition)
Part Number 377710-001

© Copyright 2004 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft®, MS Windows®, Windows®, and Windows NT® are U.S. registered marks of Microsoft Corporation.

HP 24-Port 4x Fabric Copper Switch User Guide

November 2004 (First Edition)

Part Number 377710-001

Contents

Intended Audience	ix
Typographical Conventions	ix
Contact Information	x

1: Introducing the InfiniBand System 1

About the HP 24-Port 4x Fabric Copper Switch User Guide	1
Maximize Server Networks	1
What is InfiniBand?	2
How Does InfiniBand Work?	2
Possible Components	2
Protocols	2
Architectural Elements	3
Understanding the Subnet Manager (SM)	4
Understanding Subnet Manager Routing	5

2: Getting Started 9

Overview of Set-up Procedures	9
Install HCAs and Drivers in Hosts	10
Install and Power on the InfiniBand Chassis	10
Configure Basic Connectivity	10
Configuring an Ethernet Management IP Address	10
Configuring a Direct Serial-Console Connection	11
Configuring an InfiniBand Management IP Address	12
Configuring the System Hostname	12
Test Network Connectivity	13
Verify Communication Between Hosts	13
Verify Switch to Host Communication	14
Install the GUI (Element Manager)	14
Access a Management System	14
Default User Name and Passwords	14
CLI Management	14
GUI Management	15
SNMP Management	15
(Optional) Enable Database Sync	15
Configure Partitions	15
Create a Partition ID (P_Key)	16
Specify Partition Members and the Membership Type	17
Maintain Partition Key Information	17
Set User Levels and Passwords	18

Change Default User Name and Password	18
3: Understanding the Management Options	19
About the CLI	19
Understanding the Command Modes	20
Using the CLI.....	21
Entering the Sub-Command Mode	21
Exiting Command Modes	21
Using Command Completion	21
Displaying Command History	22
Setting Terminal Parameters.....	22
Ending A CLI Session	23
Quick Help.....	24
About Element Manager.....	24
The Chassis Window	25
The Tool Bar.....	25
About Selecting Items.....	26
Using Element Manager	27
Installing the Element Manager Program	27
Starting the Element Manager	29
Reading the Element Manager Status Colors	30
About SNMP.....	30
Supported MIBs.....	30
Using SNMP	30
Configuring SNMP Settings	30
4: Performing Admin Tasks Through the GUI.....	33
Configuring the IB Interface Speed	33
Explicitly Configure IB Interface Speed	33
Set IB Interface Speed to Auto-Negotiate	34
View the IB Interface Speed.....	35
Setting the System Clock.....	36
Setting Time Manually	36
Synchronize the Clock to an NTP Server	36
Rebooting the System	37
Reboot a System with a Single Controller Card.....	37
5: Performing Admin Tasks Through the CLI.....	39
Setting the IB Interface Speed	39
Explicitly Configure IB Interface Speed	39
Set IB Interface Speed to Auto-Negotiate	40
View the IB Interface Speed.....	40
Notifying Users.....	40
Broadcasting Messages to all Users.....	41
Sending Messages to Individual Users	41

Setting the System Clock	41
Setting Time	42
Synchronize the Clock to an NTP Server	42
Rebooting the System	43
Reboot a System with a Single Controller	43

6: Setting Access and Security 45

Understanding Access and Accounts	45
About User Accounts	45
Elements of the Access System	46
Understanding Usernames and Passwords	46
About Roles and Privileges	46
Managing Access and Accounts	47
Setting or Changing a Password	47
Displaying User Information	48
Adding New Users	49
Deleting a User Account	50
User Account Configuration Commands	50
Switching User Identity	52
Changing Privilege Access-Levels	52
About Partitions	53
How Partitions Work	53
Partition Members	54
Membership Types	54
Selecting a P_Key Value	54
Understanding how P_Keys are Saved	56
Create Partitions (CLI)	56
Create a Partition ID (P_Key)	57
Specify Partition Members and the Membership Type	57
Create Partitions (GUI)	57
Create a Partition ID (P_Key)	57
Specify Partition Members and the Membership Type	58
About SSH	59

7: Using the Subnet Manager Through the GUI 61

The Subnet Manager (SM)	61
Master Subnet Manager	61
Standby Subnet Manager	62
Viewing the Subnet Manager Configurations	62
View a Summary of Subnet Management	62
View Details of Subnet Management	62
Changing the Subnet Manager Configurations	64
Change the Priority of a SM	64
Change the Sweep Interval of a SM	64

Change the Response Timeout of a SM.....	65
Managing Synchronization Between SMs.....	66
Enable/Disable Database Synchronization	66
Set Configurations for the Master SM	67
Set Configurations for the Backup SM.....	68
Adding a Subnet Manager	70
Viewing Partitions	71
About InfiniBand Multicast Groups	72
Viewing Multicast Groups.....	72
View a Multicast Group Summary	72
View Multicast Group Details	73
View the Subnet Manager Services.....	75
View a Summary of the SM Services.....	75
View Details of the SM Services.....	76
Configure Subnet Manager Routing.....	77
Configure the LID Mask Control (LMC)	78
View InfiniBand Paths.....	78

8: Using the Subnet Manager Through the CLI 81

The Subnet Manager (SM).....	81
Master Subnet Manager	81
Standby Subnet Manager	82
Viewing the Subnet Manager Configurations.....	82
View a Summary of Subnet Management	82
View Details of Subnet Management	82
Changing the Subnet Manager Configurations.....	83
Change the Priority of a SM	83
Change the Sweep Interval of a SM	84
Change the Response Timeout of a SM.....	84
Managing Synchronization Between SMs.....	84
Enable/Disable Database Synchronization	84
Set Configurations for the Master SM	85
Set Configurations for the Backup SM.....	85
Adding a Subnet Manager	87
About InfiniBand Multicast Groups	87
Viewing Multicast Groups.....	88
View a Multicast Group Summary	88
View Multicast Group Details	89
Viewing the SM Services	90
View a Summary of the SM Services.....	90
Configure Subnet Manager Routing.....	90
Configure the LID Mask Control (LMC)	91
View InfiniBand Paths.....	91

9: Using Image Files 93

Types of Image Upgrades.....	93
------------------------------	----

TopspinOS Upgrades	93
About the System Image	93
What is a System Image?	93
What is an Image File?	94
About Copying/Downloading the Image	94
Card Status Requirements	95
Upgrade Procedure Overview	95
Set-Up the Hardware Connection	95
Out-of-Band Connection	95
In-Band Connection	95
Verify the Installed Image Version	96
Check the Image Version Through the GUI	96
Check the Image Version Through the CLI	96
Copy/Download the Image	96
Copy/Download the Image Through the GUI	97
Copy/Download an Image Through the CLI	98
Activate an Image	100
Specify a New Boot Image	101
Specify a New Boot Image Through the GUI	101
Specify a New Boot Image Through the CLI	102
Reboot the System	102
Deleting Image Files	103
Deleting Images Through the GUI	103
Deleting Images Through the CLI	103
10: Using Configuration Files	105
Understanding Configuration Files	105
About the Startup-Config	105
About the Running-Config	105
Listing Configuration Files	106
List Config Files Through the CLI	106
List Config Files Through the GUI	106
Export a Configuration File	106
Export a Config File Through the CLI	107
Export a Config File Through the GUI	107
Import a Configuration File	108
Download a Config File Through the CLI	108
11: Using Log Files	111
Understanding Log Files	111
File Management and Storage	111
About Message Types	111
Listing Current Log File Names	112
Listing Current Logs Through the CLI	112
Listing Current Logs Through the GUI	112
Viewing a Log File Through the CLI	113

Display Entire Log	113
Show Most Recent Log Entries	113
Show Details of a Specific Log	114
Viewing a Log File Through the GUI	114
Filtering Logs	115
Configuring Remote Logging	117

12: Viewing the IB Network Through the GUI 119

About the Device Manager (DM)	119
Display the Device Manager	119
View I/O Unit Information	119
View I/O Controller Units	120
View I/O Controller Units Services	121
About the Topology View	121
Display the InfiniBand Topology	122
View the Topology	122
View the Name of an HCA	123
View the GUID of an HCA	124
Determine Which HCA Port is Connected to an IB Port	124
View the GUID of an IB Switch	126
Add an Attached Device to the Topology View	126
View the Internal Chassis Topology	127
View Subnet Manager Details	129
View Basic Node Information	129
View Advanced Node Information	130
View Basic Port Information	131
View Advanced Port Information	133

13: Monitoring and Reporting Through the GUI 137

About Analyzing Network Data	137
Benefits	137
Data Captured	138
About Tabular Formats	138
About Graph Formats	138
Types of Graphs	138
Creating a Data Analysis Table	140
Create a Data Table	140
Export a Data Table	141
Print a Data Table	142
Creating a Data Analysis Graph	143
Modify a Graph	145
Print a Graph	146
About SNMP Traps	146
Events Sent to Trap Receivers	146
Configuring SNMP Settings	147
Viewing Current SNMP Trap Receivers	147

Adding an SNMP Trap Receivers	147
Editing a Current SNMP Trap Receiver	148
14: Monitoring Through the CLI	149
About InfiniBand Events	149
About Tracing	149
Types of Traces	150
Trace Levels	150
About SNMP Traps	151
Events Sent to Trap Receivers	151
Configuring SNMP Settings	152
Viewing Current SNMP Trap Receivers	152
Add an SNMP Trap Receiver	152

Preface

This document is a guide to the HP 24-Port 4x Fabric Copper Switch.

Intended Audience

The intended audience is the administrator responsible for installing, configuring, and managing your equipment. This administrator should have experience administering similar networking or storage equipment.

Typographical Conventions

The following typographic conventions are used in this manual to provide visual clues as to the purpose or application of specific text.

- **Bold text** indicates a command.
- **Courier text** indicates example text as displayed on the computer screen or that you enter exactly as shown.
- **Italics** indicate variable text that you replace with an actual value.
- **Square angle-brackets** ([data]) indicate an option that you choose to include or exclude. (Do not include the brackets when supplying optional data.)
- **Piping character** (|) indicates an “or” choice. For example, a | b indicates “a or b”. [a] | [b] indicates an optional choice between a or b.
- **Menu1->Menu2->Item...** indicates a pop-up menu sequence to open a form or execute a desired function.
- **Ellipses (...)** indicate truncated text. You will see these in long examples depicting terminal output that is too long to be shown in its entirety.



NOTE: Indicates an important point or aspect that you need to consider before continuing.

Contact Information

Table 2-1: Customer Contact Information

For the name of your nearest authorized HP reseller:	In the United States, call 1-800-345-1518. In Canada, call 1-800-263-5868.
For HP technical support:	In the United States and Canada, call 1-800-HP-INVENT (1-800-474-6836). This service is available 24 hours a day, 7 days a week. For continuous quality improvement, calls may be recorded or monitored. Outside the United States and Canada, refer to www.hp.com

Introducing the InfiniBand System

This chapter gives an overview of the following:

- [“About the HP 24-Port 4x Fabric Copper Switch User Guide” on page 1](#)
- [“Maximize Server Networks” on page 1](#)
- [“What is InfiniBand?” on page 2](#)
- [“How Does InfiniBand Work?” on page 2](#)

About the HP 24-Port 4x Fabric Copper Switch User Guide

The *HP 24-Port 4x Fabric Copper Switch User Guide* is specifically intended to demonstrate the processes involved in using and managing the InfiniBand™ switch technology.

- For information regarding the Host Channel Adapter, refer to the *HP Dual-port 4x Fabric Adapter User Guide*.
- For information regarding the switch, refer to the *HP 24-Port 4x Fabric Copper Switch Hardware User Guide*.

Maximize Server Networks

The Topspin system uses InfiniBand as the underlying fabric that creates a scalable and efficient server area network. The system also seamlessly interconnects with existing Fibre Channel and Ethernet resources, extending the value of InfiniBand to the rest of the network.

What is InfiniBand?

InfiniBand (IB) is a high speed, high density serial interconnect that increases CPU utilization, decreases latency, and eases the management pain of data centers.

The term “InfiniBand” refers to the entire hardware, communication, and management infrastructure. Use of this technology increases the communication speed between:

- CPUs
- devices within servers
- subsystems located throughout a network.

How Does InfiniBand Work?

InfiniBand combines high-speed hardware, specialized protocols, and Remote Data Memory Access (RDMA) techniques to achieve the objective of increased CPU utilization and decreased latency.

Operations of the InfiniBand Architecture are managed by the Subnet Manager.

Possible Components

One or more of the following hardware components may be used to maximize your server network.

- InfiniBand switch
- Host Channel Adapters (installed in host)
- Ethernet Gateway
- Fibre Channel Gateway

Protocols

InfiniBand requires a new set of protocols. For information on how to configure these protocols, refer to the *HP Dual-port 4x Fabric Adapter User Guide*.

IPoIB

The IP over IB (IPoIB) link driver provides standardized Internet Protocol encapsulation over InfiniBand fabrics. IPoIB can transparently use IP over InfiniBand technology, similar to the way that IP runs over Ethernet.

The primary responsibilities of the IPoIB driver are to perform address resolution and the management of multicast membership.

SDP

The Sockets Direct Protocol (SDP) is a transparent protocol used on InfiniBand networks to allow sockets-based applications to take advantage of the RDMA performance over an InfiniBand network.

SDP provides:

- a reduction in the amount of software running inside a process context
- zero copy

SDP protocol support enables databases, application servers, and CPUs to operate more efficiently because the databases spend less time waiting for work, the application servers spend less time waiting for responses, and the CPUs have more cycles free for other work.

SRP

SCSI RDMA Protocol (SRP) is an upper-layer storage protocol for InfiniBand. It runs SCSI commands across RDMA-capable networks for InfiniBand hosts to communicate with Fibre Channel storage

devices. This protocol allows InfiniBand hosts to natively send SCSI commands as if the storage was direct attached.

The SRP protocol is designed to operate using an RDMA communication service. An RDMA communication service provides communication between pairs of consumers; it uses messages for control information and RDMA operations for data transfers.

The SRP protocol is only used if you have a Fibre Channel Gateway installed in your InfiniBand system.

uDAPL

The user Direct Access Programming Library (uDAPL) is a standardized user mode API that natively supports InfiniBand fabrics.

uDAPL performs name to address translations, establishes connections, and transfers data reliably.

The primary responsibilities of uDAPL are:

- Connection management
- Low latency data transfer and completion

MPI

The MPI protocol is bundled with the Upper Layer Protocol (ULP) suite. Topspin has taken the Ohio State University's (OSU's) MVAPICH and created Topspin's version of this release. However, in addition, the HCAs also run using other popular InfiniBand MPI implementations.

Alternative MPI Implementations

Topspin customers have also deployed a variety of MPIs that use Mellanox's VAPI layer. This includes OSU, LAM-MPI, Verari Systems Software, Inc's MPI/Pro (formerly Softech's), and LANL MPI. Topspin products have also been used successfully with SCALI MPI, which is based on uDAPL.

Differences Between Topspin and Standard MPI

There are significant differences between the version of MPI provided, and OSU's MPI.

- There is no restriction on which HCA port is used (OSU only supports Port 1)
- Support for Opteron 64 bit operation is provided
- Bug fixes have been provided for the purpose of improving stability

Architectural Elements

What is RDMA?

InfiniBand utilizes Remote Direct Memory Access (RDMA) technology. RDMA is a technology that allows one computer to place information directly into the memory of another computer.

RDMA is specifically characterized by two important features:

- allows user space applications to directly access hardware
- zero-copy data movement

A combination of hardware and software allows user space applications to read and write the memory of a remote system without kernel intervention or unnecessary data copies. This results in lower CPU utilization per I/O operation and more efficient use of machine resources because applications place most of the messaging burden upon InfiniBand's high-speed network hardware.

Work Queues and Queue Pairs

A "verb" is the abstract description that is used to define the functionality of the Host Channel Adapter (HCA). A "verb consumer" refers to the direct user of the verb.

A work queue provides a verb consumer with the ability to queue up a set of instructions that are executed by the Channel Adapter. There are two types of Work Queues: Send Work Queue (outbound) and a Receive Work Queue (inbound). Together these Work Queues create a Queue Pair.

The Queue Pair (QP) is one of the primary architectural elements of InfiniBand. In InfiniBand, communication occurs between Queue Pairs, instead of between ports.

A Queue Pair (QP) is an addressable entity, and consists of two Work Queues: 1). Send Work Queue and a 2). Receive Work Queue. (A work queue provides a verb consumer with the ability to queue up a set of instructions that are executed by the Channel Adapter.) The Channel Adapter hardware takes over the task of arbitrating communication - multiplexing access to the send queue or de-multiplexing messages on the receive queue.

A connection is made by linking a local queue pair to a remote queue pair. Applications do not share queue pairs; therefore, once you set them up, you can manage them at the application level without incurring the overhead of system calls.

Send and Receive work queues are:

- always created as a pair
- always remain a pair
- known as Queue Pairs
- identified by a Queue Pair number, which is within the Channel Adapter.

Queue pairs have:

- a region of memory to be used as buffers (numbers of Queue Pairs are only limited by memory).
- a key that must match on each incoming packet (the Q_Key) to verify the validity of the packet
- (potentially) a partition key, which specifies the portion of the fabric that this queue pair may access.

The queue pair is the mechanism by which you define quality of service, system protection, error detection and response, and allowable services.

Types of Services

Each queue pair is independently configured for a particular type of service. These service types provide different levels of service and different error-recovery characteristics.

The available transport-service types include:

- Reliable connection
- Unreliable connection
- Reliable Datagram
- Unreliable Datagram

Once the fabric connections are discovered, queue pairs and protection domains are established, and the type and quality of service are defined for each queue pair, the fabric operates reliably and securely at full performance without impact on system hardware or software resources.

Understanding the Subnet Manager (SM)

The Subnet Manager configures and maintains fabric operations. There can be multiple Subnet Managers, but only one master.

For information regarding configuring the subnet managers, refer to [“Using the Subnet Manager Through the GUI” on page 61](#) or [“Using the Subnet Manager Through the CLI” on page 81](#).

The Subnet Manager is the central repository of all information that is required to setup and bring up the InfiniBand fabric.

The master Subnet Manager

- Discovers the fabric topology.
- Discovers endnodes.
- Configures switches and end nodes with their parameters, such as:

- Local Identifiers (LIDs)
- Global Unique Identifier (GUIDs)
- Partition Key (P_Keys)
- Configures switch forwarding tables.
- Receives traps from Subnet Management Agents (SMAs).
- Sweeps the subnet, discovering topology changes and managing changes as nodes are added and deleted.

Understanding the Subnet Management Agents (SMAs)

Subnet Management Agents (SMA) are part of the Subnet Manager. A SMA is provided with each node and process packets from the Subnet Manager.

If an Subnet Manager is elected master, all of its components, including SA, are implicitly elected master. If a Subnet Manager ceases to be master, all of its components cease responding to messages from clients.

Subnet Manager Hot Standby

The master and slave subnet managers can be synchronized so the information in the master is carried over to the slave in the event of a fail-over. Refer to [“Enable/Disable Database Synchronization” on page 84](#) to configure SM hot standby.

The hot standby/database sync feature is used to synchronize the databases between subnet managers running on separate chassis.

The Subnet Manager maintains a data base in the volatile memory of the master SM containing all required information.

How is the synchronization done?

The database synchronization is accomplished in two stages:

- Cold Synchronization - This stage is initiated by the master SM when it is ready to start a synchronization session with a standby SM. In this stage, all out of sync tables are copied from the master SM to the standby SM.
- Transactional Synchronization - This stage is entered following successful completion of the cold synchronization stage. In this stage, all database update transaction requests that are processed by the master, are replicated to the standby.

What can cause a standby SM to become the master SM?

- A crash of the node running the current master SM.
- Partitioning of the subnet (e.g. due to link failure).
- Graceful shutdown of the master (e.g. for maintenance purposes).

What happens when a master subnet manager fails?

In the event of a failure:

- The standby subnet manager becomes the new master.
- The new master rebuilds the data base from information retrieved during the subnet discovery phase.
- Existing LID assignments are retained, where possible.
- All ports are reset to force them to re-join multicast groups, re-advertise services, re-request event forwarding, and re-establish connections.
- A “SlaveToMaster” event trap is generated to trigger any necessary processing by external management applications.

Understanding Subnet Manager Routing

There are two different concepts associated with InfiniBand routing:

- Routing internally within a switch (hops between switch chips)
- Routing between whole switches (hops between nodes). This is also referred to as routing between “switch elements.”

Internal switch routing can be configured to provide the highest performance in passing traffic, and to minimize the threat of congestion within the switch.

The Routing Process Overview

1. The Subnet Manager (SM) first discovers all the InfiniBand switch chips in the network.
2. The SM groups the internal switch chips within each chassis into a “switch element.”
3. The SM process continues until all the InfiniBand switches are grouped into “switch elements.”
4. After all the switch chips are grouped, the SM will route the switch elements according to the routing algorithm discussed in [“Minimum Contention, Shortest Path & Load Balancing Algorithm” on page 6.](#)
5. The internal network of each InfiniBand switch is then routed based on the best algorithm for each “switch element.”

Multiple Paths

The SM allows you to define the Logical Identifier Mask Control (LMC) value per subnet. The default value of the LMC is 0, so by default only one Logical Identifier (LID) is assigned to each host port.

Once the LMC value has been assigned, the SM will route different paths for each LID associated with the same host port. The result of these paths is based on the routing algorithm applied.

Understanding SM Routing Terms

The following terms are important to understand before distinguishing the various types of algorithms that the Subnet Manager uses for routing:

Distance - Distance is defined as the number of hops (InfiniBand switches or “switch elements”) between source and destination.

Contention - A contention is declared for every switch port on the path that is already used for routing another LID associated with the same host port.

Minimum Contention, Shortest Path & Load Balancing Algorithm

Minimum Contention, Shortest Path and Load Balancing is the algorithm that is used by default to route between the “switch elements” and for routing between the internal InfiniBand switch chips *within* each “switch element.”

The following algorithm is used for the calculation:

1. The shortest path for each of the host ports is calculated.
2. Contention is calculated for all the available paths that are within the (shortest path + tolerance) distance.
 - a. The path with the least contention is selected.
 - b. If two paths have the same contention, the path with less distance is selected.
 - c. If two paths have the same contention and the same distance, the port usage count is used to provide load balancing over the two paths. The usage count is a measure of how many LIDs have been configured to use that particular port.

Configuring Your Network For Optimal Routing

Create Equal Paths Between Switch Elements

It is recommended that InfiniBand switch elements be connected so that all paths between any pair of switch elements are the same distance (i.e. same number of hops), if possible. This enables you to obtain the optimal paths using the default tolerance of 0.

Determine the First Path that will be Discovered

The SM Routing Algorithm selects the first best path that it finds. If multiple paths with the same properties are available then the first of these paths found is the one that is selected. Therefore, it is possible to setup the cabling between switch elements to force the algorithm to prioritize certain paths. Depending on the network requirements, the prioritized paths can either be concentrated on a particular switch element or spread across multiple switch elements to improve fault-tolerance.

Getting Started

The information in this chapter focuses on the software and firmware aspects of the initial set-up, and assumes that you have additional documentation for the hardware.

This chapter provides the following information:

- Overview of entire system installation on [page 9](#), with references to more detailed information.
- Setup procedures for the InfiniBand™ switch.

Overview of Set-up Procedures

Follow the steps below to configure the InfiniBand server switch system.

1. Determine your hardware topology.
2. Install the Host Channel Adapter and drivers ([page 10](#)).
3. Install and power-on the InfiniBand Chassis ([page 10](#)).
4. Configure Basic Connectivity ([page 10](#)).
5. Test Network Connectivity ([page 13](#)).
6. Install the Element Manager GUI ([page 14](#)).
7. Access a Management System ([page 14](#)).
8. Configure Partitions ([page 15](#)).
9. Set User Level and Access ([page 18](#)).

Install HCAs and Drivers in Hosts

Refer to the *HP Dual-port 4x Fabric Adapter Quick Setup Installation* card and the *HP Dual-port 4x Fabric Adapter User Guide*.

Install and Power on the InfiniBand Chassis

Refer to the *HP 24-Port 4x Fabric Copper Switch Hardware Quick Setup Installation* card and the *HP 24-Port 4x Fabric Copper Switch Hardware User Guide* for installation and power instructions.

Configure Basic Connectivity

The InfiniBand switch is not pre-configured with an IP address. You must configure the IP address of a management port to administer and monitor the InfiniBand switch with the CLI and Element Manager. A Management port is provided for a connected Ethernet host running TCP/IP or connected InfiniBand hosts running IPoIB. Configure the Management port you wish to use.

- [“Configuring an Ethernet Management IP Address” on page 10](#)
- [“Configuring an InfiniBand Management IP Address” on page 12](#)

Configuring an Ethernet Management IP Address

To configure an out-of-band Ethernet Management IP address:

1. Make sure that the InfiniBand switch is attached to a PC or terminal via the serial port. Refer to the *HP 24-Port 4x Fabric Copper Switch Hardware Quick Setup Installation* card and the *HP 24-Port 4x Fabric Copper Switch Hardware User Guide*.
2. Open a terminal emulation program, such as HyperTerminal for Windows®, and set the session parameters as follows:
 - Baud: 9600 b/s
 - Data Bits: 8
 - Parity: None
 - Stop Bits: 1
 - Flow control: None
3. At the Login: prompt, enter the username and password. The default is **super** and **super**.

Example

```
Login: super
Password: super
Topspin-360>
```

4. At the CLI prompt, enter **enable**. This enters the privileged-execute mode.

```
Topspin-360> enable
```

5. Enter **configure** to enter the global-configuration mode.

```
Topspin-360# configure
Topspin-360(config)#
```

- Set the IP address and netmask. The following address is an example.

```
Topspin-360 (config) # interface mgmt-ethernet
Topspin-360 (config mgmt-ethernet) # ip address 10.10.0.22 255.255.255.0
```

- Set the default gateway address. This address is an example.

```
Topspin-360 (config mgmt-ethernet) # gateway 10.10.0.1
```

- Enable the management port

```
Topspin-360 (config mgmt-ethernet) # no shutdown
```

- Test IP connectivity by pinging the management station.

```
Topspin-360 (config mgmt-ethernet) # exit all
Topspin-360 > ping 10.10.0.3
sending 5 ICMP Echos to 10.10.0.3, 56 data bytes
!!!!

Success rate is 100 percent (5/5)
round-trip min/avg/max = 0.000000/0.000000/0.000000 ms
Topspin-360 >
```

- You must save the configuration persistently by using the **copy** command.

```
Topspin-360 (config mgmt-ethernet) # exit
Topspin-360 # copy running-config startup-config
```

You are now ready to power down the chassis and mount it. Later, you can configure the box via Telnet, SSH, Chassis Manager or Element Manager.

Configuring a Direct Serial-Console Connection

Refer to the *HP Serial Management Cable Guide* for information regarding setting up the physical serial-console connection.

Remote Telnet Login

You can Telnet to the Management-Ethernet port on the box from a host on the same network as the Management-Ethernet port, or from any host with a route to the Management-Ethernet network.

To run the CLI remotely:

- Open a terminal or terminal emulator window.
For example: from the command line, enter the telnet command with the IP address, or network name, of the Management-Ethernet port.

```
# telnet 10.0.0.47
```

The CLI login prompt (Login:) is displayed.

- Enter a CLI user name.
The CLI password prompt (Password:) is displayed.

- Enter the CLI user password.
The prompt changes to indicate a successful login. The system is now ready to receive CLI commands.

Remote SSH Login

TopspinOS supports SSH2 for secure, encrypted login to the CLI. SSH is enabled by default, and does not require additional configuration.

To login via SSH:

1. Use an SSH client (e.g. Putty) to port 22.

Configuring an InfiniBand Management IP Address

To configure an In-band InfiniBand management IP address:

1. At the Login: prompt, enter the username and password. The default is **super** and **super**.

```
Login: super
Password: super
Topspin-360>
```

2. At the CLI prompt, enter **enable**. This enters the privileged-execute mode, as indicated by the # sign.

```
Topspin-360> enable
Topspin-360#
```

3. Enter **configure** to enter the configuration mode.

```
Topspin-360# configure
Topspin-360(config)#
```

4. Enter the interface to be configured, and set the IP and mask addresses.

```
Topspin-360(config)# interface mgmt-ib
Topspin-360(config mgmt-ib)# ip address 10.3.102.20 255.255.255.0
```

5. Set the default gateway address. The gateway address refers to the address of the internal port.

```
Topspin-360(config mgmt-ib)# gateway 10.3.0.1
```

6. Enable the IB management port.

```
Topspin-360(config mgmt-ib)# no shutdown
```

7. Test IP connectivity by pinging an InfiniBand host on the other side of the gateway.

```
Topspin-360(config mgmt-ib)# exit all
Topspin-360> ping 10.3.102.34
sending 5 ICMP Echos to 10.3.102.34, 56 data bytes
!!!!

Success rate is 100 percent (5/5)
round-trip min/avg/max = 0.000000/0.000000/0.000000 ms
Topspin-360>
```

8. Save the configuration by using the **copy** command, or wait until you execute the **reload** command. You will be prompted to save the unsaved configuration changes. .

```
Topspin-360(config mgmt-ethernet)# exit
Topspin-360# copy running-config startup-config
```

You are now ready to power down the chassis and mount it. Later, you can configure the box via Telnet, SSH, or the Element Manager.

Configuring the System Hostname

The Topspin system allows you to assign a hostname to the system for management purposes.

To assign a hostname name to the management port:

1. Start a CLI session.
2. Enter the privileged-user mode.

```
Topspin-360> enable
Topspin-360#
```

3. Enter the global-configuration mode.

```
Topspin-360# configure
Topspin-360 (config) #
```

4. Enter the **hostname** command with the name you wish to assign.

The **hostname** command assigns a convenient name to the system that shows up at the CLI prompt.

```
Topspin-360 (config) # hostname MyHost
```



NOTE: This command also changes the CLI prompt. The new hostname is applied immediately, however, the prompt does not change until you change modes. For example, the prompt changes when you exit the global-configuration mode.

Test Network Connectivity

Refer to the *HP 24-Port 4x Fabric Copper Switch Hardware User Guide* for information regarding connecting network devices.

After you install network cables, you can verify connectivity by pinging those connected devices from the CLI or pinging between attached hosts.

Verify Communication Between Hosts

To verify the device recognizes and successfully links InfiniBand-attached hosts, enter the **ping** command on one host and specify the IP address of another connected host.

```
# ping 10.2.65.50
PING 10.2.0.50 (10.2.0.50) from 10.2.0.41 : 56(84) bytes of data.
64 bytes from 10.2.0.50: icmp_seq=0 ttl=64 time=164 usec
64 bytes from 10.2.0.50: icmp_seq=1 ttl=64 time=144 usec
...
...
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max/mdev = 0.135/0.147/0.164/0.017 ms
#
```

Verify Switch to Host Communication

To verify the InfiniBand device can reach a host on the network, enter either the user-execute or privileged-execute mode on the InfiniBand device, then enter the `ping` command. This is an In-band procedure only.

```
Topspin-360# ping 10.10.253.47
Sending 5 ICMP Echos to 10.10.253.47, 56 data bytes
!!!!!!

Success rate is 100 percent (5/5)
round-trip min/avg/max = 0.000000/0.000000/0.000000 ms
Topspin-360#
```

Install the GUI (Element Manager)

HP 24-Port 4x Fabric Copper Switches can be managed visually through the Element Manager, which provides a wide range of configuration, monitoring, and troubleshooting options.

Refer to [“Installing the Element Manager Program” on page 27](#) for information regarding the Element Manager installation.

Access a Management System

Default User Name and Passwords

For initial configuration, log in as the unrestricted user.

- The default unrestricted username for the CLI is **super** and the default password is **super**.
- The default community-string assigned to this user for the Element Manager is **secret**.

Use the following methods to manage the Topspin system.

CLI Management

Refer to [“About the CLI” on page 19](#) for more information about managing through the CLI.

Run the Command Line Interface (CLI) from one of the following methods:

- [“Direct Serial-Console Connection” on page 14](#)
- [“Remote Telnet Login” on page 14](#)
- [“Remote SSH Login” on page 15](#)

Direct Serial-Console Connection

Refer to the *HP Serial Management Cable Guide* for information regarding setting up the physical serial-console connection.

Remote Telnet Login

You can Telnet to the Management-Ethernet port on the box from a host on the same network as the Management-Ethernet port, or from any host with a route to the Management-Ethernet network.

To run the CLI remotely:

1. Open a terminal or terminal emulator window. For example: from the command line, enter the `telnet` command with the IP address, or network name, of the Management-Ethernet port.

```
# telnet 10.0.0.47
```

The CLI login prompt (`login:`) is displayed.

2. Enter a CLI user name.

The CLI password prompt (`password:`) is displayed.

3. Enter the CLI user password.

The prompt changes to indicate a successful login. The HP 24-Port 4x Fabric Copper Switch system is now ready to receive CLI commands.

Remote SSH Login

TopspinOS supports SSH2 for secure, encrypted login to the switch CLI. SSH is enabled by default, and does not require additional configuration.

1. To login via SSH, use an SSH client (e.g. Putty) to port 22.

GUI Management

1. Refer to [“About Element Manager” on page 24](#) for more information about managing through the CLI.
2. Run the Element Manager (GUI) over a TCP/IP network.
3. To log in to the GUI, refer to [“Starting the Element Manager” on page 29](#).

SNMP Management

For more information regarding SNMP, refer to [“About SNMP” on page 30](#).

Any Network Manager running the Simple Network Management Protocol (SNMP) can manage the Topspin system, if the Management Information Base (MIB) is installed correctly. By default, the Topspin GUI is a network manager and uses SNMP v2c as the protocol to communicate between the chassis and the management workstation.

(Optional) Enable Database Sync

If you are configuring more than one InfiniBand chassis in your fabric, it is likely that you will want to enable database synchronization of the subnet managers.

- To enable data synchronization with the Element Manager GUI, refer to [“Enable/Disable Database Synchronization” on page 66](#).
- To enable data synchronization with the CLI, refer to [“Enable/Disable Database Synchronization” on page 84](#).

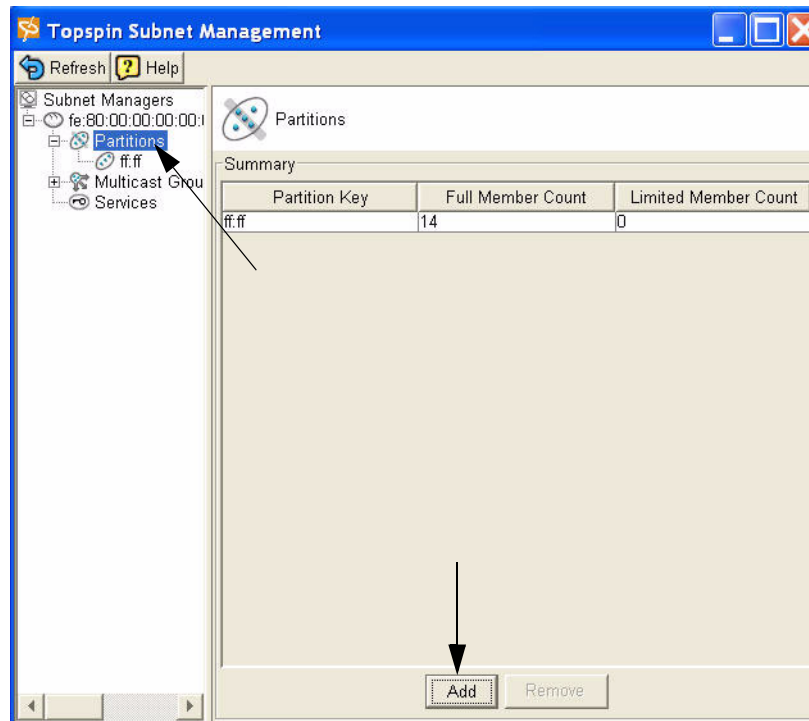
Configure Partitions

Partitions are described in detail in [“About Partitions” on page 53](#).

Create a Partition ID (P_Key)

A default partition is configured automatically. The members of a default partition include all connected ports, and provide full membership. However, to create separation between traffic, you must configure specific partitions.


1. Launch Element Manager, if you have not already done so.
2. Select **InfiniBand --> Subnet Management**.
The Subnet Management window appears.
3. Click open the **Subnet Manager** folders in the left window.
The Partitions folder appears.
4. Click on the **Partitions** folder in the left window. The Partitions Summary window appears.



5. Click the **Add** button.

The Add Partition dialog box appears.

Enter a Partition key (P_Key) to identify the new partition. For information regarding selecting values, refer to the “[Selecting a P_Key Value](#)” on page 54.




- a. Click the **Add** button.

The new Partition appears in the left window.

Specify Partition Members and the Membership Type

- b. Click on the new Partition in the left window.

The available partition members appear in the right-side window.



Note that the “Available Members” refers only to members that are known to the Subnet Manager. This includes HCAs and switches that are already plugged into the fabric as well as manually configured entries.

If you know the GUID and port count of an HCA that has not yet been installed, you can configure it before it is plugged in by using the “Add Other” button.

- c. Click on a member from the Available Member list, and use the arrow button to move it to the Full or Limited member columns.

For information regarding Membership Types, refer to the [“Membership Types” on page 54](#)

- d. Click back to the Partitions folder (in the left-side window) when you have selected all of the members for your Partition.

The new Partition appears in the Partition Summary window.

Maintain Partition Key Information

The configured p_keys will be needed in completing the configuration of the system.

- e. Configured partition keys must be mapped to any of the following components that exist:
- Host Channel Adapters (HCAs). Refer to the *HP Dual-port 4x Fabric Adapter User Guide*.
 - Ethernet Gateway Bridge-groups.

- Fibre Channel gateways.
- f. If you have multiple InfiniBand switches in your fabric:
 - Exchange the partition configuration between switches by enabling database synchronization, if you have not already done so. Refer to [“Enable/Disable Database Synchronization” on page 84](#).

Set User Levels and Passwords

Change Default User Name and Password

For security purposes, since multiple users exist on the system, it is highly recommended that you change the default passwords after initial configuration.

See [“Understanding Usernames and Passwords” on page 46](#) for more information.

1. Log in to the CLI as a super user. Use the default username (**super**) and the default password (**super**) if they have not already been changed (refer to [page 47](#)).
2. Enter the privileged-user mode.
3. Enter the global-configuration mode.
4. Enter the **username** command and the **password** keyword to change the user account and user password.

Use the default user name and password if they have not already been changed (refer to [page 47](#)).

The user name and password are alphanumeric strings of up to 34 characters each.

5. Repeat step 4 to change additional usernames and passwords.

Example

```
Topspin-360# Login: super
Password: xxxx
Topspin-360> enable
Topspin-360# configure
Topspin-360(config)# username ib-fc_admin password ibFcAdmin
Topspin-360(config)# username ib-fc_admin communitystring
ibFc-commStr
```

6. Exit the global-configuration mode.
7. Use the **show user** command to verify changes.

Only a user with unrestricted privileges may view user information.

```
Topspin-90> show user all
=====
User Information
=====
      username : admin
      password : topspin
      snmp-community : justatest
      permission-level : ib-rw, ip-ethernet-rw, fc-rw
      admin-status : enabled
      num-logins : 0
      num-unsuccessful-logins : 0
      last-login :
      last-unsuccessful-login :

Topspin-90>
```

Understanding the Management Options

This chapter gives an overview of the following system Management options:

The CLI

- [“About the CLI” on page 19](#)
- [“Using the CLI” on page 21](#)

The Java GUI

- [“About Element Manager” on page 24](#)
- [“Using Element Manager” on page 27](#)

The Web GUI

- Refer to the *HP 24-Port Fabric Copper Switch Chassis Manager User Guide*

SNMP

- [“About SNMP” on page 30](#)
- [“Using SNMP” on page 30](#)

About the CLI

The Topspin system can be managed through the Command Line Interface. For more information regarding the CLI, refer to the *HP 24-Port Fabric Copper Switch Command Line Reference Guide*, or [“Understanding the Command Modes” on page 20](#).

The CLI includes the following features:

- IOS-like syntax
- Command Completion
- Context Help

- Multiple Command Modes

Example

```
# telnet topspin_90
Login: super
Password: xxxx
Topspin-90> enable
Topspin-90#
```

Understanding the Command Modes

The CLI has four command modes

- user-execute mode (read-only)
- privileged-execute mode
- global-configuration mode
- sub-command mode

The commands you can enter depend upon the current command mode and who you log in as. You may enter a question mark (?) at the CLI prompt to list the commands appropriate for the current mode and user identity.

User-Executive Mode

The user-execute mode is the entry point to the privileged-execute mode and all CLI sessions begin in the user-execute mode. This mode provides commands for viewing some of the HP 24-Port 4x Fabric Copper Switch configuration and some user information. Guest users may only work in the user-execute mode.

Privileged-Execute Mode

The privileged-execute mode can view the entire switch configuration and all user information. It is used to perform some high-level administrative tasks, such as saving the current configuration and setting the system clock. It is also the access point to the global-configuration and sub-command modes. You must enter the privileged-execute mode before entering the configuration modes.

Use the **enable** keyword to enter the privileged-execute mode. Note that only administrative and unrestricted users may enter the privileged-execute mode.

```
# telnet topspin_90
Login: super
Password: xxxx
Topspin-360> enable
Topspin-360#
```

Mode changes are reflected in changes to the Topspin system prompt.

For example, going from the user-execute to privileged-execute mode, the prompt changes from Topspin-90> to Topspin-90#.

Global-Configuration Mode

Enter the global-configuration mode from the privileged-execute mode. The global-configuration mode is used to configure everything except interface cards and their ports. The global-configuration mode configures system-level attributes, such as SNMP, SNMP agents, and the networks.

Enter the `config` keyword while in the privileged-execute mode to enter the global-execute mode.

```
Topspin-90# configure
Topspin-90(config)#
```


Sub-Command Mode

The final mode is sub-command mode. Anything to do with InfiniBand, Ethernet, and Fibre Channel interface cards, ports, and gateways is done in this mode, including the Management-Ethernet ports. This mode is used to assign IP addresses to interface gateway ports, set connection speeds, set connection types, etc.

Using the CLI

Entering the Sub-Command Mode

1. Enter global-configuration mode
2. Enter the interface keyword
3. Enter the type of interface to be configured

For example, to enter the interface-configuration mode for configuring the Management-Ethernet port, enter:

```
Topspin-90 (config) # interface mgmt-ethernet
Topspin-90 (config-if-mgmt-ethernet) #
```

Exiting Command Modes

Most commands are mode-dependent. For example, you can only log out of a Topspin system session in the user-execute or privileged-execute mode. To configure the Topspin system, you will have to enter and exit Topspin system modes.

The `exit` command is used to return to the previous mode.

```
Topspin-360 (config-if-fc-5/1) # exit
Topspin-360 (config) # exit
Topspin-360 #
```

You may also enter the `exit` command with the `all` argument to return to the user-execute mode in one step. If you are currently in the privileged-execute mode, `exit` with the `all` keyword will log you out of the CLI session.

To exit the privileged-execute mode and return to the user-execute mode, enter the `disable` command. For example:

```
Topspin-360 (config) # disable
Topspin-360 >
```

Using Command Completion

The system provides command completion by way of the [Tab] key. If you enter a partial command and press the [Tab], the CLI will complete the command and place the cursor at the end of the command.

To facilitate command entry, CLI commands do not have to be entered in their entirety. You may enter just enough of each command or argument to make it uniquely identifiable.

For example:

```

Topspin-360 (config) # fc ?
  srp                - Configure FC SRP
  srp-global         - Configure FC SRP-global parameters
Topspin-360 (config) # fc srp- ?
  enable            - Enable FC SRP
  gateway-portmask-pol - Configure FC SRP-global gateway-portmask-policy
  itl               - Configure FC SRP-global ITL
  lun-policy        - Configure FC SRP-global lun-policy
  target-portmask-poli - Configure FC SRP-global target portmask policy
Topspin-360 (config) # fc srp- gate ?
  restricted        - Configure FC SRP gateway-portmask-policy restricted
Topspin-360 (config) # fc srp- gate res ?
<cr>
Topspin-360 (config) # fc srp- gate res

```

In the preceding example, `srp-` is short for `srp-global`, `gate` is short for `gateway-portmask-pol`, and `res` is short for `restricted`.

Note: Command completion only works for commands; it is not effective for keywords.

Displaying Command History

The Topspin system “remembers” the last 40 commands you enter.

Display the commands in the command history by using the following command:

```
history
```

You can also use up and down arrows to toggle between commands.

Setting Terminal Parameters

The TopspinOS can be customized to set the number of lines displayed at one time by commands like **more** (used to prevent data from scrolling quickly out of view). The number of lines specified only applies to the current CLI session. Other users are unaffected by changes to the display length.

You can also set a limit for inactivity during a CLI session. Changes to this parameter are applied immediately to all users, whether logged in now or who log in later.

1. Enter the **length** command

```

terminal length int
terminal no length

```

length int is a number between 1 and 512 that indicates the number of lines to display at one time. It is recommended you set the terminal page length to 0 when using the `tail` command with the `end` argument. Otherwise, you will have to keep pressing the `<space>` bar to continue each time the maximum display length is reached.

`no` resets the terminal length to the default (24 lines per page).

```

Topspin-360# terminal time-out 45
Topspin-360# terminal no time-out

```

time-out int is a number between 1 and 100000 that indicates the number of minutes of inactivity allowed before automatically closing the CLI session.

`no` resets the session limit for inactivity to the default (15 minutes).

Example

To set the number of lines displayed on the terminal screen to 67 lines at a time and raise the time-out limit to 60 minutes

1. Enter the user-execute or privileged-execute mode.
2. Enter the `terminal` command with the `length` parameter and the number of lines to display per page.

```
Topspin-90# terminal length 67
```

3. Enter the **terminal** command with the **time-out** parameter and the number of minutes to allow a user to remain inactive before closing their CLI session.

```
Topspin-90# terminal time-out 60
```

4. Verify the changes were made by entering the **show terminal** command.

```
Topspin-90# show terminal
Console is enabled
Connection host address is 10.10.253.47
Length: 67 lines, Width: 80 columns
Timeouts: enabled, Value: 60 minutes
Session limit is set to 3
History is enabled, history size is 30
Maximum command length is 512 characters
Maximum login attempts is 5
Topspin-90#
```

Ending A CLI Session

To terminate the current CLI session:

1. Enter `exit all` to return to the user-execute mode.
2. You may want to save the current configuration so that it is re instantiated the next time the system reboots.

```
Topspin-360# copy system:running-config config:startup-config
```

The **copy** command “copies” the current configuration to the `startup-config` file, which is used to reconfigure the chassis upon reboot. The `running-config` is a virtual file that contains all unsaved configuration commands. When it is “copied” the system saves the `running-config` into the `startup-config` file so that they can be maintained across reboots.

3. Enter the CLI **logout** command.

```
Topspin-360# logout
Topspin-360# Login:
```

Users who initiated a CLI session through a Telnet or SSH connection to the Management-Ethernet port will be logged out and the connection closed. Users connected directly to the Serial-Console port will still be prompted to login to the CLI.

NOTE: Do not use `<control> c` to terminate an active CLI session. It has been disabled to ensure that the CLI is always running on the terminal. Only the CLI `copy` command recognizes `<control> c`.

Quick Help

You can enter the question mark (?) at the CLI prompt to display one of three types of user information.

1. Enter a question mark (?) at the CLI prompt at any time to display the commands you can execute. Only the commands appropriate to the current mode and user login are displayed.
2. You may also enter part of a command string, and a question mark, to display the choices for completing that command string.

You may enter just enough of each command or argument to make it uniquely identifiable, followed by a space and a question mark to display the arguments and keywords you need to continue the command line.

You may enter just enough of each command or argument to make it uniquely identifiable, followed by a space and a question mark to display the arguments and keywords you need to continue the command line. For example:

```

Topspin-360 (config) # i? <-- display all keywords that start with "i"
Configure Commands:
  ib          - IB subnet manager configuration
  ib-agent    - Configure IB agent settings
  interface   - Select an interface to configure
  ip          - Global IP configuration subcommands
Topspin-360 (config) # in? <-- display all keywords that start with "in"
interface    <-- only 1 keyword starts with "in"
Topspin-360 (config) # in ? <-- display the arguments to "interface"
ethernet    - Configure Ethernet interfaces
fc          - Configure Fibre Channel interfaces
gateway     - Configure Gateway settings
ib          - Configure InfiniBand interfaces
mgmt-ethernet - Configure Ethernet Management port
mgmt-ib     - Configure InfiniBand Management port
Topspin-360 (config) # in <-- waits for you to complete the "interface"
                    command string

```

After displaying the help information, the system enters the command string up to the question mark on the input line and waits for you to complete the string. You do not have to retype the string.

If there is no space between a partially-entered command string and the question mark, a list of possible completions are displayed, as shown above.

When enough characters have been entered to uniquely identify a command or keyword in a command string, you may leave it as-is, enter a space, and then add additional keywords or arguments, or you can press the <Tab> key to complete the commands or keywords to improve readability.

About Element Manager

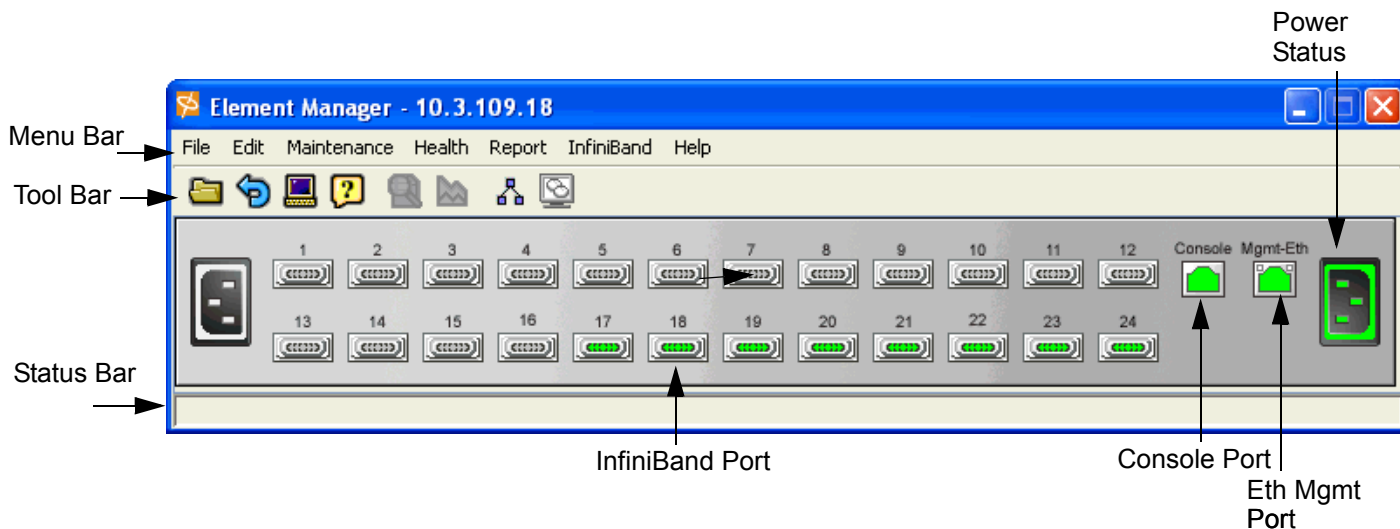
The Element Manager is the Graphic User Interface that can be used to manage most of the Topspin system functionality.

The Element Manager is comprised of several different areas, which allows you to manage the entire Topspin system.

The Chassis Window

Once you log into the Element Manager, the chassis window is displayed. This is the primary window in which you work. It graphically depicts the current configuration of the attached InfiniBand system chassis.

Sample HP 24-Port 4x Fabric Copper Switch Chassis View



The Tool Bar

The tool bar contains a set of icons for commonly used functions. These are described in the table below.

Table 3-1: Element Manager Tool Bar










	Initiates a Topspin system session on a different chassis. The new window opens and requests the host address and user credentials. It is the equivalent to selecting File->Open...
	Polls the physical chassis for current connectivity and status, and then updates the display. It is the equivalent to selecting File->Refresh .
	Opens a terminal window and starts a Telnet session with the physical switch chassis that is hosting the current session. It prompts for CLI user credentials. Once you supply this information, you may use CLI commands to configure and monitor the switch chassis. It is equivalent to selecting File->Telnet . See “Remote Telnet Login” on page 11 .
	Displays context-sensitive usage information about the current window. It is the equivalent to selecting Help->Contents . See “Quick Help” on page 24 .

Table 3-1: Element Manager Tool Bar

	<p><i>Port Configuration:</i> Opens a window that displays the configuration of the selected port(s). For most ports, the window is also use to change port configuration. It is equivalent to selecting Edit->Port Properties....</p> <p><i>Card Configuration:</i> If the selected object is a card, the Edit window opens. There is no pull-down menu equivalent for this function when a card is selected; use Edit-> Card Properties.... .</p> <p>or</p> <p>Double-click the card or right-click the card and select Edit from the pop-up menu.</p>
	<p>Opens a window that displays selected statistical data in a line chart format. It is typically used to chart and analyze network traffic statistics. There is no specific menu equivalent; however this icon is typically included on most windows opened by Report --> Graph->Port.... or Report --> Graph->Card.... Refer to “About Analyzing Network Data” on page 137.</p>
	<p>Opens the Topology view, a graphic tool that depicts the switch and channel adapter connections of the current InfiniBand fabric configuration. Refer to “About the Topology View” on page 121.</p>
	<p>Opens the Subnet management window. Configure and view current Subnet Managers, Partitions, and Multicast groups. It is the equivalent to selecting InfiniBand -> Subnet Manager.</p>
	<p>Opens the Storage Manager window. Configure and view current Fibre Channel SRP information. It is the equivalent to selecting Fibre Channel -> Storage Manager.</p>

About Selecting Items

Interaction with the Topspin system is performed using a combination of pull-down menus, icons in the tool bar, and pop-up menus. The windows these open contain a combination of text fields, radio buttons, and toggles with which to configure selected cards and/or ports.

Depending upon the function you wish to execute, you select one or more cards and/or ports in the following ways:

- by placing the cursor over the desire object(s) and clicking the left mouse button.
- by placing the mouse cursor over the object and right-clicking the mouse. This displays a pull-down menu from which you may select an item.
- by placing the mouse cursor over it and double-click the left mouse button. This selects the object and opens a default window. The window displayed depends upon the selected object. This is explained in more detail in the sections that follow.

Everything in the switch chassis display may be selected. You can select the chassis, interface cards, ports, management ports, and serial console ports. The only exception are the cards without ports. These are placeholders for later expansion and cannot be selected.

Select multiple objects on the switch display by using <Control> left-click. However, when you select an object that is not the same type as the currently selected set, the selected set is de-selected.

For example, if you <Control> left-click multiple Ethernet ports and then attempted to select a Fibre Channel port:

The Ethernet ports are de-selected and the Fibre Channel port is the only thing selected.

Using Element Manager

Installing the Element Manager Program

For information regarding upgrading the Element Manager, refer to [“Starting the Element Manager” on page 29](#).

HP 24-Port 4x Fabric Copper Switch devices can be managed visually through the Element Manager, which provides a wide range of configuration, monitoring, and troubleshooting options.

The Element Manager runs on multiple platforms, including Windows NT/2000/XP and Linux. To install the Element Manager:

1. Check that you have sufficient system resources.

You will need:

- 64 MBytes free RAM
- 75 MBytes disk space + 50MBytes additional temporary space during installation
- 300 Mhz processor
- 800x600 screen resolution, 16bit color depth

2. Locate software.

- a. Go to <http://support.hp.com/>
- b. Select “Software & Driver downloads.”
- c. On the Software & Driver Downloads page, enter your product name, then click the double arrow.

3. Install the software.

- a. Unzip the tar file containing the software using gunzip.
- b. Extract the software into a local directory using tar.
- c. Change directory to Element Manager (EM)

4. Locate the correct install program in the Linux or Windows directory for your architecture. Execute that program.

For example, for Linux ia32:

- `cd Linux`
- `./install_linux_x86.bin`

5. Click **Next**.

The Choose Install Folder window opens.

6. Specify the folder in which to install the software. You may:

- Enter the full path to where you want the software installed.

or

Click the **Choose** button to browse for a folder.

- Click **Restore Default Folder** to restore the folder location to its original value. On Windows, the default is `C:\Program Files\Topspin Element Manager`. On Solaris and

Linux, the default is `/home/TopspinEM`, where *home* is the home directory of the person installing the software.

If the folder does not exist, you will be prompted to create it.

7. Click the **Next** button.

On Windows, the Choose Shortcut Folder window opens.

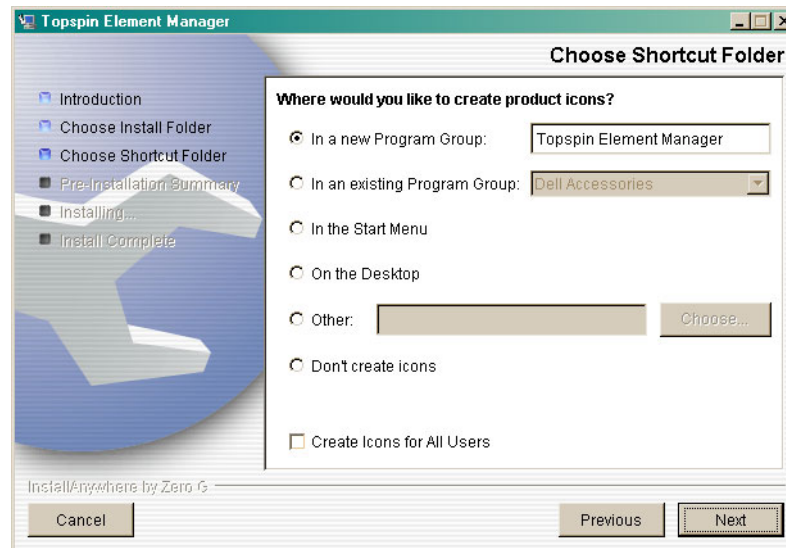


Figure 3-1: Element Manager Installation, Choose Shortcut Folder Window

On Linux, the Choose Link Folder opens.

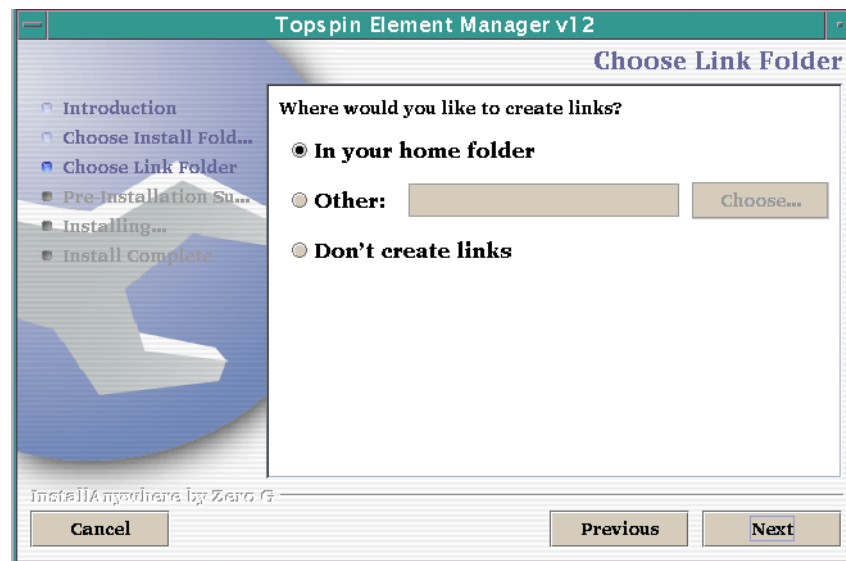


Figure 3-2: Element Manager Installation, Choose Link Folder Window

8. Specify where you want shortcuts, or links, to the Element Manager placed.

You may select multiple options. You may also specify a unique placement in the Other field or by clicking the **Choose...** button.

If you want to change the settings in a preceding window, click **Previous**.

9. Click the **Next** button.

The Pre-Installation Summary window is opened. This window lists the installation choices you have made thus far.

10. If you are satisfied with your installation choices, click Install.

The Installing Element Manager window opens to indicate installation status.

If you are not satisfied with the configuration, click Previous to return to the preceding window and make the desired changes.

The Installation Complete window opens when the software is installed.

11. Click Done.

The window closes and Element Manager installation is complete.

Starting the Element Manager

1. To start the Element Manager:

On a Windows system, select the Element Manager icon or select it from the Programs menu. For example:

Start->Programs->Topspin Element Manager->TopspinEM

The **Open Device** window opens. This window is used to specify the IP address or DNS name of the management port and the administrator's community string. The community string functions as an SNMP password.

If you are on a UNIX system, change to the directory containing the Element Manager executable, or add the directory to your search path. Enter the Element Manager command on the command line.

Example

```
# ./TopspinEM
```

The **Open Device** window opens.

2. Enter the IP address or network name of the management port in the **Device Name** field. Enter the IP address or network name of the out-of-band management port or the in-band management port.




Figure 3-3: Element Manager, Open Device window

The IP address of the management port in the [Figure 3-3](#) example is 10.10.253.47.

3. Enter the appropriate community string in the **SNMP Community** field.

The default unrestricted community string is **secret**. For information regarding community strings, refer to [“Configuring SNMP Settings” on page 30](#).

4. Click **Open**.

A graphic representation of the switch chassis and current configuration is displayed.

The Element Manager is now ready to configure the InfiniBand network, as well as Ethernet or Fibre Channel expansion module(s).

The Element Manager is dynamically updated to show changes to the configuration. As cards and ports are configured, the corresponding run lights and port frames reflect the changes by turning green. Depending upon your Element Manager Preference settings, it may take a few seconds for configuration changes to be shown in the Element Manager display.

Reading the Element Manager Status Colors

The colors in the Element Manager display indicate the state of each port:

Table 3-2: Interpreting the Element Manager Port Colors

Color	Indication
Green	A link is established between the port and a connected host or switch. A host or switch must be connected to the port for it to turn green. A green port does not indicate network activity.
Grey	The port is enabled but there is no link, as in the case of a disconnected cable.
Red	The port is disabled.
Transparent	The port is unmanaged.

About SNMP

The Topspin system can also be managed via SNMP v2C, including a variety of MIBs and Traps.

Any Network Manager running the Simple Network Management Protocol (SNMP) can manage the switch, if the Management Information Base (MIB) is installed correctly. By default, the switch GUI is a network manager and uses SNMP as the protocol to communicate between the chassis and the management workstation.

Supported MIBs

In addition to private MIBs, the Topspin system supports the standards MIBs listed below.

- RFC2665: Ethernet-Like MIB
- RFC1213: MIB2
- RFC2863: The Interface Group
- RFC2096: IP Forward MIB
- IB SM - Draft InfiniBand Subnet Manager
- IB SMA - Draft InfiniBand Subnet Management Agent

Using SNMP

Configuring SNMP Settings

The following SNMP parameters can be configured on the system:

- **Authorized Trap Receivers** - You can set one or more network management stations on your network to receive traps. By default, the Element Manager is configured to be an authorized trap

receiver. You can have a maximum of six trap receivers. Entries in this list can be configured from the **Health --> Trap Receivers** menu.

- **Community Strings** - You can set community strings as a simple method of authentication between the system and the remote Network Manager. One unique community string is associated with each username and password. Community strings can be associated with a variety of privilege levels. For a list of default community strings associated with each user, see [“Managing Access and Accounts” on page 47](#).

Performing Admin Tasks Through the GUI

This chapter describes the following administrative procedures that can be performed through the Element Manager GUI:

- [“Configuring the IB Interface Speed” on page 33](#)
- [“Setting the System Clock” on page 36](#)
- [“Rebooting the System” on page 37](#)

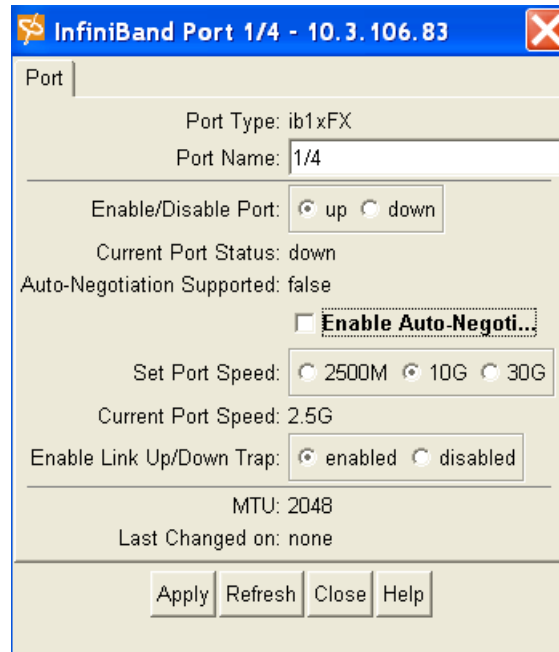
Configuring the IB Interface Speed

Explicitly Configure IB Interface Speed

To explicitly set the speed for the InfiniBand interface ports as 1x or 4x:

1. Confirm that you are using the appropriate InfiniBand cable for the speed you intend to set.
For example, if you intend to set the speed as 4x, it is imperative that you confirm you are using a 4x cable. Using a 1x cable on a 4x speed InfiniBand interface will cause serious performance issues.
2. Launch Element Manager, if it is not already open.
A graphic representation of your InfiniBand switch appears.
3. Double-click an InfiniBand port.

The IB Port window appears.



4. Uncheck the **Enable Auto-Negotiate** box. Leave auto-negotiate checked if you want the speed of the transmit port and the receiving port to automatically negotiate the highest possible speed.
5. Select the 2500M or 10G speed. The 30G speed is not yet supported.
6. Click the **Apply** button.

Set IB Interface Speed to Auto-Negotiate

Set the InfiniBand interface speed to auto-negotiate if you want the speed of the transmit port and the receiving port to automatically negotiate the highest possible speed.

1. Launch Element Manager, if it is not already open.
A graphic representation of your InfiniBand switch appears.
2. Double-click an InfiniBand port.

The IB Port window appears.

3. Check the **Enable Auto-Negotiate** box.
4. Click the **Apply** button.

View the IB Interface Speed

View the current speed of the InfiniBand interfaces:

1. Launch Element Manager, if it is not already open.
A graphic representation of your InfiniBand switch appears.
2. Double-click an InfiniBand port.

The IB Port window appears.

3. Check the **Enable Auto-Negotiate** box.

4. Click the **Apply** button.

Setting the System Clock

Maintaining accurate time is important for statistics and auditing. The InfiniBand chassis provides an on-board system clock to save the time settings across reboots. Time is maintained in one of two ways:

- An on-board system clock
- External NTP servers

Time is set at the factory, and can be manually set. To ensure accurate synchronization, it is recommended that you use an external NTP server. This enables logs to be synchronized with other management systems.

Setting Time Manually

To set the system clock in the Element Manager:

1. In the Element Manager, select **Maintenance** -> **Time...**
2. Enter the time manually in the **Time** and **Date** fields.

It is recommended that you use NTP servers to maintain the system clock because it assures accuracy and avoids the potential time drift inherent to unsynchronized clocks.

3. Click **Apply**.
4. Click **Close**.

Synchronize the Clock to an NTP Server

This section shows you how to use the Element Manager to synchronize the Topspin system clock with an NTP server:

1. In the Element Manager, select **Maintenance** -> **Time...**

The **Chassis Time** window opens.

2. Ignore the **Date** and **Time** group; it will be reset automatically by the NTP servers.
3. In the **NTP Servers** group, enter the IP addresses of the NTP servers in the **NTP Server1**, **NTP Server2**, and **NTP Server3** fields.


4. Click **Apply**.
5. Click **Close**.

Rebooting the System

Reboot a System with a Single Controller Card


1. Launch the Element Manager.
2. Choose the reboot option:
 - a. Right-click in any unused part of the Element Manager display, including unused interface slots.

A pop-up window opens that allows you to select **Reboot**.




- b. or select **Maintenance** --> **Reboot**.

If changes have been made, you will be asked if you want to save the changes:



3. Select **Yes** to save changes, or **No** to discard changes.

A prompt appears to verify your desire to reboot the system.



4. Click **Yes** to reboot the system or **No** to return to the Element Manager display.

Upon rebooting, the connection is lost or the Element Manager opens a "Timeout" dialog. These indicate the reboot process is taking place.

Performing Admin Tasks Through the CLI

This chapter describes the following administrative procedures that can be performed through the CLI.

- [“Setting the IB Interface Speed” on page 39](#)
- [“Notifying Users” on page 40](#)
- [“Setting the System Clock” on page 41](#)
- [“Rebooting the System” on page 43](#)

Setting the IB Interface Speed

InfiniBand interface port speeds can be configured to 1x or 4x.

Explicitly Configure IB Interface Speed

To explicitly set the speed for the InfiniBand interface ports as 1x or 4x:

1. Confirm that you are using the appropriate InfiniBand cable for the speed you intend to set.
For example, if you intend to set the speed as 4x, it is imperative that you confirm you are using a 4x cable. Using a 1x cable on a 4x speed InfiniBand interface will cause serious performance issues.
2. (Optional) Disable auto-negotiate if it is currently enabled. The interface speed cannot be configured as long as auto-negotiate is enabled.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config) interface ib all no auto-negotiate
Topspin-360(config-if-ib-1/1-1/12) #
```

3. Enter the following command:

```
config interface ib [{IB switch card/IB switch port | all } speed 1x | 4x ]
```

Example

```
Topspin-360> enable
Topspin-360# config interface ib all speed 4x
```

Set IB Interface Speed to Auto-Negotiate

Set the InfiniBand interface speed to auto-negotiate if you want the speed of the transmit port and the receiving port to automatically negotiate the highest possible speed. IB speed can be set on an individual port basis, or all at once.

1. Enter the following command:

```
config interface ib [{IB switch card/IB switch port | all } auto-negotiate]
```

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# interface ib all
Topspin-360(config-if-ib-1/1-1/12)# auto-negotiate
Topspin-360(config-if-ib-1/1-1/12)#
```

View the IB Interface Speed

View the current speed of the InfiniBand interfaces:

1. Enter the following command:

```
show interface ib [IB switch card/IB switch port | all]
```

Example

```
Topspin-360> enable
Topspin-360# show interface ib 15/1
=====
                        InfiniBand Interface Information
=====
                        ....
                        ....
auto-negotiate-supported : yes
      auto-negotiate     : enabled
      admin-speed        : 10gbps
      oper-speed         : unknown
                        ....
                        ....
*****
```

Notifying Users

User notification commands send text messages to the terminal screens of all CLI users or to individual users. These are convenient utilities for notifying everyone of an impending reboot or to notify single users about special issues that apply only to them.

Broadcasting Messages to all Users

Message broadcasting is an important feature to forewarn all CLI users that some major event is about to take place, such as bringing down a network for administration. A broadcast message is sent to every active CLI session on the InfiniBand system chassis.

Enclose multi-word messages within double-quotes. Single-word messages do not require double-quotes. Only the unrestricted user may broadcast messages.

Syntax

```
Topspin-90# broadcast "message"
```

The message identifies the sender, followed by the message text.

For example, if you send this:

```
Topspin-90# broadcast "FC card 5 going down in 10 minutes."
```

Everyone, including the user who sent the message, receives this:

```
Topspin-90# Broadcast message from super
```

```
FC card 5 going down in 10 minutes.
```

Sending Messages to Individual Users

The **write** command is used to send a message to a single user. Check that the user is logged in before attempting to write to their terminal.

1. Enter the **show user user_name** command to verify the user is logged in.
2. Enter the **write** command to message the individual.

```
Topspin-90> show user waldo
=====
                        User Information
=====
      username : waldo
      access-level : readwrite
      admin-status : enabled
      status : active
      num-logins : 1
      num-unsuccessful-logins : 6
      last-login : Thu Oct 10 09:13:10 2002
      last-unsuccessful-login : Thu Oct 10 09:12:32 2002
Topspin-90> write waldo "Connection to FC array 15 is now
working."
Topspin-90>
```

Setting the System Clock

Maintaining accurate time is important for statistics and auditing. The switch chassis provides an on-board system clock to save the time settings across reboots. Time is maintained in one of two ways:

- An on-board system clock
- External NTP servers

Time is set at the factory, and can be manually set. To ensure accurate synchronization, it is recommended that you use an external NTP server. This enables logs to be synchronized with other management systems.

Setting Time

Note: If you have an NTP server configured, it is recommended that you use the process described in: [“Synchronize the Clock to an NTP Server” on page 42.](#)

To set the Topspin system clock in the CLI:

1. Login to the CLI as the `super` user.
2. Enter `enable` to enter the privileged-execute mode.

```
Topspin-360> enable
Topspin-360#
```

3. Enter the `clock` command, followed by the `set` keyword and the time and date in the format `hh:mm:ss dd mm yy`. For example:

```
Topspin-360# clock set 19:22:10 25 03 03
```

4. Save your configuration.

```
Topspin-360# copy running-config startup-config
```

Synchronize the Clock to an NTP Server

You can set the InfiniBand switch to synchronize the time with up to three NTP servers.

To set the InfiniBand system clock in the CLI:

1. Log in to the CLI as the `super` user.
2. Enter `enable` to enter the privileged-execute mode.

```
Topspin-360> enable
Topspin-360#
```

3. Enter the `ntp` command, and the keyword `server-one` before entering the IP address. This is the first server to which the IB switch will synchronize.

```
Topspin-360# config
Topspin-360(config)# ntp server-one 10.0.3.10
Topspin-360(config)#
```

4. Enter the IP address of a second NTP server.

```
Topspin-360(config)# ntp server-two 10.0.3.11
Topspin-360(config)#
```

5. Enter the IP address of a third NTP server.

```
Topspin-360(config)# ntp server-three 10.0.3.13
```

Rebooting the System

Reboot a System with a Single Controller

Enter the CLI **reload** command in privileged EXEC mode. The system prompts for you to verify the reload. If you had not already saved configuration changes, and the system detects the changes, it prompts you to save. If you enter yes, the new configuration is stored in **startup-config**. You may optionally save the configuration to an alternate file by entering the file name, followed by a carriage-return.

Example

```
Topspin-360> enable
Topspin-360# reload
System configuration has been modified. Save?
[yes(default)/no/*.cfg] yes
Proceed with reload? [confirm]
Topspin-360#
Connection to host lost.
```

The system re-initializes itself and then loads the system-image and the startup-config file. Wait a few minutes and attempt to log onto the chassis.

Setting Access and Security

This chapter describes the following Access and Security features:

- [“Understanding Access and Accounts” on page 45](#)
- [“Managing Access and Accounts” on page 47](#)
- [“About Partitions” on page 53](#)
- [“Create Partitions \(CLI\)” on page 56](#)
or
- [“Create Partitions \(GUI\)” on page 57](#)
- [“About SSH” on page 59](#)

Understanding Access and Accounts

About User Accounts

A user account is used to control who gains access to the Topspin system. Access can be achieved through the CLI (console, telnet, SSH) and SNMP. CLI access is authorized through a password validation. SNMP access is authorized through a community-string validation.

User accounts can be added, deleted, and modified as needed. Up to 15 user accounts are supported. Only user(s) that have the unrestricted ReadWrite permission level can add, delete, and modify user accounts. Each Topspin system is preconfigured with three factory default user accounts.

Each user account can be administratively enabled and disabled as needed. The user can disable a user account without having to delete it from the system.

Each user account is uniquely identified by an ascii string that can be up to 20 characters long. No two user accounts can have the same user name.

In order for users to initiate an administration session, the User has to supply login credentials. The credentials supplied depend upon the interface being used.

Elements of the Access System

The Topspin access system associates the following key elements:

- Username - Creates a unique username in the system.
- Password
- Community String - Unique string used for the Element Manager and SNMP Network Managers.
- Privilege Level - Allows combinations of different privileges.

The CLI uses username and password. The Element Manager uses the community string to identify which user has logged in. Granular access rights are given by privilege level.

Understanding Usernames and Passwords

CLI users enter standard username and password information to begin a CLI session. By default, you may log on as one of three users, `super`, `admin`, or `guest`. The user names are shown in the table below.

Table 6-1: Default User Names, Passwords, and Privileges

User Name	Default Passwords	Privileges
super	By default, the password is “super”. The default community string is “secret”.	The super user has unrestricted privileges. Use this account to manage any part of the Topspin system. This user may view and modify a configuration, as well as administer user accounts and access privileges. This user configures the console and management ports for initial chassis setup.
admin	By default, the password is “admin”. The default community string is “private”.	The admin user has general read-write privileges. This user may view and modify the current configuration. However, the admin user can change only its own user information, such as the admin password.
guest (disabled by default)	The default password is “guest”. The default community string is “public.”	The guest user has read-only privileges. This user may only view the current configuration. The guest user cannot make any changes during the CLI session.

About Roles and Privileges

Roles allow granular levels of privileges. For example, you can create separate Fibre Channel, Ethernet, or InfiniBand administrators, who only have access to specific subsystems. The Topspin system combines multiple roles with read and read-write access for flexible control. These roles are enforced with both the CLI and the Element Manager.

The unrestricted administrator (`super`) is responsible for assigning these roles. Network administrators are given read-only and read-write access to each of the three network types.

Understanding Permission Levels

The following table displays the different access-levels.

Table 6-2: Description of Privilege Levels

Level	Description
ib-ro	InfiniBand read-only access.
ib-rw	InfiniBand read-write access.
ip-ethernet-ro	Ethernet read-only access.
ip-ethernet-rw	Ethernet read-write access.
fc-ro	Fibre Channel read-only access.
fc-rw	Fibre Channel read-write access.
unrestricted-rw	Read-write access to all network configuration commands.

Managing Access and Accounts

Setting or Changing a Password

1. Log in to the CLI as a super user. Use the default username and password if they have not already been changed (refer to [page 14](#)).
2. Enter the privileged-user mode.
3. Enter the global-configuration mode.
4. Enter the **username** command and the **password** keyword to change the user account and user password.

Use the default user name and password if they have not already been changed (refer to [page 14](#)).

The user name and password are alphanumeric strings of up to 34 characters each.

5. Repeat step 4 to change additional usernames and passwords.

Example

```
Topspin-360# Login: super
Password: xxxx
Topspin-360> enable
Topspin-360# configure
Topspin-360(config)# username ib-fc_admin password ibFcAdmin
Topspin-360(config)# username ib-fc_admin communitystring
ibFc-commStr
```

6. Exit the global-configuration mode.

7. Use the **show user** command to verify changes.

```

Topspin-360# show user
=====
User Information
=====
username : testuser
password : $1$OHJt61CE$ANK02CcPqKnFoxJ0AKAtB.
snmp-community : secret
permission-level : unrestricted-rw
admin-status : enabled
num-logins : 4
num-unsuccessful-logins : 0
last-login : Mon Nov 17 22:43:25 2003
last-unsuccessful-login :
Topspin-360#

```

Displaying User Information

To display the information of configured users:

1. Log in to the CLI as a super user.
Only a user with unrestricted privileges may view user information.
2. Enter the **show user all** command to display current user information.

```

Topspin-90> show user all
=====
===
User Information
=====
===
username : admin
password : topspin
snmp-community : justatest
permission-level : ib-rw, ip-ethernet-rw, fc-rw
admin-status : enabled
num-logins : 0
num-unsuccessful-logins : 0
last-login :
last-unsuccessful-login :
...
...
...
username : super
password : super
snmp-community : secret
permission-level : unrestricted-rw
admin-status : enabled
num-logins : 1
num-unsuccessful-logins : 0
last-login : Tue Nov 19 10:03:47 2002
last-unsuccessful-login :
Topspin-90>

```

Adding New Users

To add a new user account for both CLI and Element Manager access:

1. Log in as the unrestricted user.

Only a user with unrestricted permissions may add new user accounts.

```
Topspin-360# Login: super
Password: xxxx
Topspin-360>
```

2. Enter the privileged-user mode.

```
Topspin-360> enable
Topspin-360#
```

3. Enter the global-configuration mode.

```
Topspin-360# configure
Topspin-360 (config) #
```

4. Create the user account and user password.

The user name and password are alphanumeric strings up to 34 characters each.

```
Topspin-360 (config) # username ib-admin password ib-passwd
```

where *ib-admin* is the name assigned to this user account, *password* is a mandatory keyword, and *ib-passwd* is the password for this user account.

5. Assign an SNMP community string to the new user account.

The user must have an SNMP community string to begin an Element Manager session. If you do not want users to have SNMP access to the Topspin system, do not assign them a community string. By default, a new user account has a null or empty community string.

```
Topspin-360 (config) # username ib-admin community-string ib-commStr
```

where, *ib-admin* is the name of the user account, *community-string* is a mandatory keyword. *ib-commStr* is the SNMP community string for this user account.

6. Assign access privileges.

By default, the new user account has read-only access. You may grant write privileges to the user for functional areas, such as InfiniBand, Ethernet, and Fibre Channel.

7. Enter multiple access privileges in the order shown in [“About Roles and Privileges” on page 46](#).

```
Topspin-360 (config) # username ib-admin privilege ib-rw ip-ethernet-ro fc-rw
```

where, *ib-admin* is the name of the user account, *privilege* is a mandatory keyword, and *ib-rw*, *ip-ethernet-ro*, and *fc-rw* are access privileges. Valid access privileges are described in [“About Roles and Privileges” on page 46](#).

All new user accounts are now enabled and ready for use.

8. Exit the global-configuration mode.

- View the new user account information.

```

Topspin-360 (config) # exit
Topspin-360# show user ib-admin
=====
===
                                User Information
=====
===
      username : ib-admin
      password : ib-passwd
      snmp-community : ib-commStr
      permission-level : ib-rw, ip-ethernet-ro, fc-rw
      admin-status : enabled
      num-logins : 0
      num-unsuccessful-logins : 0
      last-login :
      last-unsuccessful-login :
Topspin-360#

```

Deleting a User Account

- Log in as the unrestricted user.
Only a user with unrestricted permissions may create and modify user accounts.

```

Topspin-360# Login: super
Password: xxxx
Topspin-360>

```

- Enter the privileged-user mode.

```

Topspin-360> enable
Topspin-360#

```

- Enter the global-configuration mode.

```

Topspin-360# configure
Topspin-360 (config) #

```

- Enter the **username** command with the name of the user, and the **no** keyword.

For example:

```

Topspin-360 (config) # username ib-admin no

```

User Account Configuration Commands

Use the following commands and keywords to administer User Accounts:Community Strings.

In the Element Manager, unique SNMP “community strings” act as user passwords. Because these are unique to each username, the community string determines which administrator is logged on. The privilege level is defined separately.

Use the following commands and keywords to administer User Accounts:

Table 6-3: User Account Administrative Commands

Command	Task
username <i>user</i> password <i>passwd</i>	Creates a new user account.
username <i>user</i> no	Deletes a user account.
username <i>user</i> community-string <i>string</i>	Assigns a community string to a user account.
username <i>user</i> no-community-string	Removes a community string from a user account.
username <i>user</i> privilege <i>priv1</i> [<i>priv2</i>] [<i>priv3</i>]	Assigns one or more permission level(s) to a user account. Refer to “Privileges” on page 46 for a list of privilege levels. Example: ib-ro, ib-rw, ip-ethernet-ro, ip-ethernet-rw, fc-ro, fc-rw, unrestricted-rw
username <i>user</i> no permission	Assigns a default permission level to a user.
username <i>user</i> enable	Administratively enables a User Account.
username <i>user</i> disable	Administratively disables a User Account.

For rapid access, the Element Manager saves the IP address and community string of recent administrators. These address/community string pairs are displayed in a scroll-list the next time you want to open an Element Manager session. Note that whoever logs on after you can reuse this connection information and, if you connected with superuser permissions, the person who follows after you shall be able to do so, too.

You can disable this functionality by referring to “Disabling the Element Manager Auto-Connection” on page 51.

Disabling the Element Manager Auto-Connection

To disable automatic connection saving in the Element Manager:

1. Select **File->Preferences...**
2. Select the **Misc** tab.
3. Uncheck the **Save communities in lastopen file** toggle.

Changing Community Strings

Use the CLI to set or change SNMP community strings. The user must have a SNMP community string to begin an Element Manager session. If you do not want users to have SNMP access to the Topspin system, remove their community string.

To change community strings:

1. Log in as the unrestricted user.

Only a user with unrestricted permissions may create and modify user accounts. However, any user with write access (administrative and unrestricted) can remove their own community string.

```
Login: super
Password: xxxx
Topspin-360>
```

2. Enter the privileged-user mode.

```
Topspin-360> enable
Topspin-360#
```

3. Enter the global-configuration mode.

```
Topspin-360# configure
Topspin-360 (config) #
```

4. Enter the `username` command with the name of the user, the `community-string` keyword, and the new community string to assign this user.

```
Topspin-360 (config) # username ib-admin community-string ib-commStr
```

In the example above, `ib-admin` is the name of the user, `community-string` is a mandatory keyword. `ib-commStr` is the SNMP community string the user will have to use to begin Element Manager sessions or execute other SNMP functions.

Switching User Identity

One of the following scenarios may make it necessary to change your user identity during a CLI session:

- you created a new user account and you want to verify the access privileges
- you have multiple administrative user-accounts and you want to switch to another administrative area

To change your user identity:

1. Enter the `user-execute` or `privileged-execute` mode.
2. Enter the `login` command with the name of a Topspin system user.

```
Topspin-90# login admin
```

3. Enter the user password.

After you enter the password, you are logged in as the specified user in the user-execute mode.

```
Password: xxxxxx
Topspin-90>
```

Changing Privilege Access-Levels

1. Log in as the unrestricted user.

Only a user with unrestricted permissions may create and modify user accounts.

```
Topspin-360# Login: super
Password: xxxx
Topspin-360>
```

2. Enter the privileged-user mode.

```
Topspin-360> enable
Topspin-360#
```

3. Enter the global-configuration mode.

```
Topspin-360# configure
Topspin-360 (config) #
```

4. Enter the `username` command with the name of the user, the `privilege` keyword, and the privileges to assign this user. For example:

```
Topspin-360 (config) # username ib-admin privilege ib-rw ip-ethernet-ro
fc-rw
```


In the example above, *ib-admin* is the name of the user account, **privilege** is a mandatory keyword, and *ib-rw*, *ip-ethernet-ro*, and *fc-rw* are access privileges.



NOTE: When you change a user's privileges, all the old privileges are removed and replaced with the new privilege(s). If the user had multiple privileges, include the other privileges on the command line when you make the change. Privileges are order-dependent. Enter them in the same order as shown in [Table 6-2 on page 47](#).

Example

The following example gives a user read-write access to InfiniBand and Ethernet configuration commands.

```

Login: super
Password: xxxx
Topspin-360> enable
Topspin-360# configure
Topspin-360(config)# username IB_admin privilege ib-rw ip-ethernet-rw
fc-ro
Topspin-360(config)# exit
Topspin-360# show user IB_admin
=====
                                User Information
=====
                                username : IB_admin
                                password : $1$LZHfWO6k$6LSXKZ7adbcC6/WXXBTAF/
                                snmp-community : IB_admin
                                permission-level : ib-rw, ip-ethernet-rw, fc-ro
                                admin-status : enabled
                                num-logins : 0
                                num-unsuccessful-logins : 0
                                last-login :
                                last-unsuccessful-login :
Topspin-360#

```

About Partitions

A partition defines a set of InfiniBand nodes that are permitted to communicate with one another. Partitions provide:

- Security
- Allows a large cluster to be divided and isolated into small “sub-clusters.”
- Maps IB nodes to selected VLANs

How Partitions Work

A partition defines a set of InfiniBand nodes that are permitted to communicate with one another. Each node may be part of multiple partitions so that a system administrator can define overlapping partitions as the situation requires. Normal data packets carry a 16-bit P_Key, or partition key, that defines a unique partition. The subnet manager configures each node's channel adapter with its set of P_Keys. When a packet arrives at a node, the channel adapter checks that the packet's P_Key is valid based on the subnet manager's configuration. Packets with invalid P_Keys are discarded. P_Key validation prevents a server from communicating with another server outside of its partition.

InfiniBand partitions are comparable to hardware-enforced security features of conventional I/O networking technologies, such as Ethernet VLANs and Fibre-Channel zones.

Partition Members

Without members, a Partition doesn't have meaning to the system. Ports are added to the Partition, and become members of that Partition. Each port may be part of multiple partitions so that the system administrator can define overlapping partitions as the situation requires.

At the time a port member is added to the Partition, the administrator must decide whether that particular port will have full or limited membership.

Membership Types

A Partition contains a group of members, but different types of members can exist within a single partition.

Partition memberships allows even further control because it defines communication within the members of that group, and not just outside of it.

There are two types of partition memberships: full membership, and limited membership

Table 6-4: Membership Types

Port Membership Types	
A Partition contains Partition Members (ports). A single Partition can contain both full or limited members.	
Full Membership	Limited Membership
A full-membership Partition Member can communicate with all other Partition Members, including other full members, as well as limited members.	A limited-membership Partition Member cannot communicate with other limited-membership partition members. However, a limited Partition Member can communicate with a full member.

About the Default Partition

The Subnet Manager automatically configures a default partition, which is always p_key ff:ff.

The default partition controls all connected ports, and by default, everything is a full-member of the default partition. The default p_key cannot be altered or deleted as it is the controlling mechanism that manages the configuration of all the partitions.

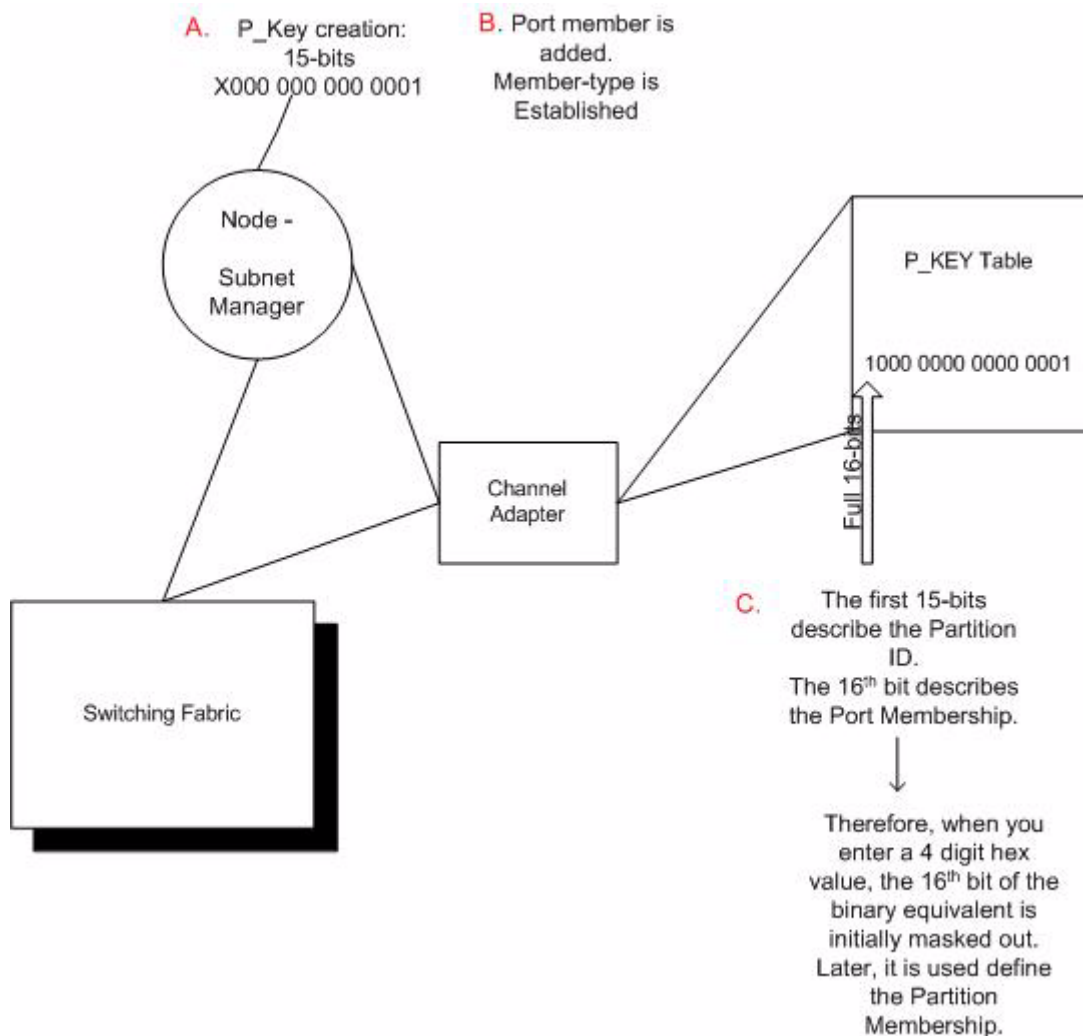
Selecting a P_Key Value

For a list of acceptable P_Key values, refer to [Table 6-6 on page 56](#).

Upon creation, the p_key value is technically a 15-bit number. However, after the p_key is created and the port(s) membership type has been established, the entire value becomes 16-bits. The most significant bit (MSB) displays the type of membership (0 = Limited member, 1 = Full member).

When assigning a p_key value, you need to choose 4 hexadecimal numbers. However, because of the way that the 16th bit is used, only certain numbers can be used for the left-most variable (the MSB). For example, do not create two p_keys:

0#:## and 8#:##, as they will be viewed as the same number by the system.



Hex to Binary Conversions

The following table is provided to assist in the creation of P_keys.

When creating the Partition p_key, enter a hexadecimal value that is the equivalent of 16-bits in binary. For example, enter 80:00 (hex) to be 1000000000000000 (binary).

The default Partition (which cannot be altered) is 7f:ff.

Table 6-5: Binary Conversions

Hexadecimal	Binary
0	0000
1	0001
2	0010
3	0011
4	0100
5	0101
6	0110

Table 6-5: Binary Conversions

Hexadecimal	Binary
7	0111
8	1000
9	1001
A	1010
B	1011
C	1100
D	1101
E	1110
F	1111

Examples of Valid P_Key Values

You can choose your own p_key values, or you can simply choose your values from the list in the following table.

Table 6-6: Valid P_Key Values

Valid P_Key Numbers	
00:01	00:11
00:02	00:12
00:03	00:13
00:04	00:14
00:05	00:15
00:06	00:16
00:07	00:17
00:08	00:18
00:09	00:19
00:10	00:20

Understanding how P_Keys are Saved

Partition information is saved persistently by the master subnet manager. P_key information can be synchronized between the master subnet manager and a slave subnet manager. The synchronization of the subnet managers means that the partition configuration (as well as other information) is exchanged between the active and standby subnet managers. Therefore, the partition configuration will be transferred in the event that an InfiniBand should fail.

The partition configuration is not saved persistently on a slave subnet manager.

If you have more than one InfiniBand switch in your fabric, refer to [“Enable/Disable Database Synchronization” on page 66](#).

If you are configuring one InfiniBand switch, it will automatically be the master, and the partition configuration is saved persistently on the switch.

Create Partitions (CLI)

Partitions are described in detail in [“About Partitions” on page 53](#).

Create a Partition ID (P_Key)

Default partitions are configured automatically. The members of a default partition include all connected ports, and provide full membership. However, to create separation between traffic, you must configure specific partitions.

Create a partition using the following steps:

- a. Enter the following items at the global-configuration mode prompt:
 - the **ib sm subnet-prefix** command
 - the subnet-prefix of the IB subnet (use the **show ib sm config subnet-prefix all** command).
 - the **p_key** command
 - and
 - an ID for the partition (refer to [page 56](#) to select a value).

```
Topspin-360 (config)# ib sm subnet-prefix 255.255.0.0 p_key 00:01
Topspin-360 (config)#
```

Specify Partition Members and the Membership Type

- b. Add the following information for partition members:
 - the **ib sm subnet-prefix** command
 - the subnet-prefix that is to be partitioned.
 - the **p_key** command
 - the current p_key value
 - the **partition-member** command
 - the GUID of the node that you want to add to the partition.
 - the port number that is to be added to the partition.
 - the membership type of the partition member (full-member or limited-member) refer to [“Membership Types” on page 54](#).

```
Topspin-360 (config)# ib sm subnet-prefix 255.255.0.0 p_key 00:01
partition-member 00:05:ad:00:00:00:02:30 1 full-member
Topspin-360 (config)# exit
```

Create Partitions (GUI)


Partitions are described in detail in [“About Partitions” on page 53](#).

Create a Partition ID (P_Key)


Default partitions are configured automatically. The members of a default partition include all connected ports, and provide full membership. However, to create separation between traffic, you must configure specific partitions.

1. Launch Element Manager, if you have not already done so.
2. Select InfiniBand --> Subnet Management.
The Subnet Management window appears.

- Click open the Subnet Manager folders in the left window.
The Partitions folder appears.
- Click on the Partitions folder in the left window. The Partitions Summary window appears.



- Click the Add button.
The Add Partition dialog box appears.
Enter a Partition key (P_Key) to identify the new partition. For information regarding selecting values, refer to the [“Examples of Valid P_Key Values”](#) on page 56.




- Click the Add button.
The new Partition appears in the left window.

Specify Partition Members and the Membership Type

- Click on the new Partition in the left window.

The available partition members appear in the right-side window.



Note that the “Available Members” refers only to members that are known to the Subnet Manager. This includes HCAs and Switches that are already plugged into the fabric as well as manually configured entries.

If you know the GUID and port count of an HCA that has not yet been installed, you can configure it before it is plugged in by using the **Add Other** button.

8. Click on a member from the Available Member list, and use the arrow button to move it to the Full or **Limited member** columns.

For information regarding Membership Types, refer to the [“Membership Types” on page 54](#).

9. Click back to the Partitions folder (in the left-side window) when you have selected all of the members for your Partition.

The new Partition appears in the Partition Summary window.

About SSH

In addition to Telnet, the CLI can be accessed via the Secure Shell (SSH2) protocol to enable a secure session. This provides strong authentication and secure communications over insecure channels. This protects the system against common security attacks, such as IP spoofing, IP source routing, and interception of clear-text passwords.

Using the Subnet Manager Through the GUI

This chapter provides the following information:

- [“The Subnet Manager \(SM\)” on page 61](#)
- [“Viewing the Subnet Manager Configurations” on page 62](#)
- [“Changing the Subnet Manager Configurations” on page 64](#)
- [“Managing Synchronization Between SMs” on page 66](#)
- [“Adding a Subnet Manager” on page 70](#)
- [“Viewing Partitions” on page 71](#)
- [“About InfiniBand Multicast Groups” on page 72](#)
- [“Viewing Multicast Groups” on page 72](#)
- [“View the Subnet Manager Services” on page 75](#)
- [“Configure Subnet Manager Routing” on page 77](#)

The Subnet Manager (SM)

The subnet manager configures and maintains fabric operations. It is the central repository of all information that is required to setup and bring up the InfiniBand fabric.

Subnet managers are identified by their subnet prefix and Global Unique Identifier (GUID).

There can be multiple subnet managers, but only one master.

Master Subnet Manager

The subnet manager that is authoritative, or has the reference configuration information for the subnet.

Standby Subnet Manager

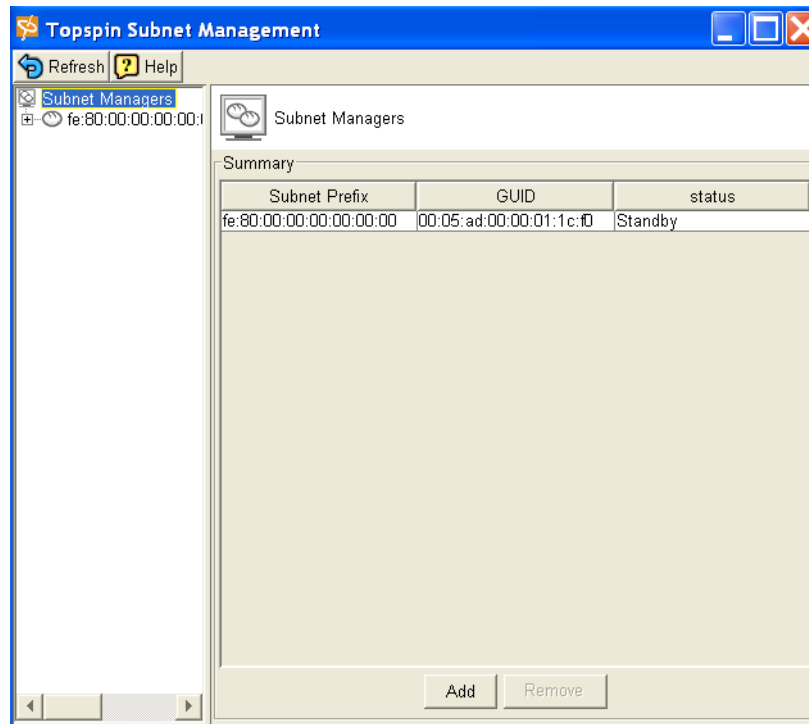
A standby subnet manager is a subnet manager (SM) that is currently quiescent, and not in the role of a master SM. Standby SMs are dormant managers, and can take over in case of failure of the master subnet manager.

Viewing the Subnet Manager Configurations

View a Summary of Subnet Management

1. Launch the Element Manager, if you have not already done so.
2. Select **InfiniBand > Subnet Management**

The Subnet Management window appears.




3. View a summary of the current subnet manager(s).
4. Continue to [“View Details of Subnet Management”](#) on page 62 to view details of the subnet management.

View Details of Subnet Management

5. Open a summary view of the subnet management. Refer to [“View a Summary of Subnet Management”](#) on page 62.
6. Click on a specific subnet manager from the left navigation bar.

Information specific to that subnet manager appears.



- View the subnet-prefix of the subnet manager.
- View the Global Unique Identifier (**GUID**) of the subnet manager.
- View the **Status** of the subnet manager.
This is the operational status, as determined by self-detection. The values are notActive, discovering, or Master. As there is only one subnet manager running on the fabric, the sm that is running is always designated the master.
 - notActive indicates the subnet manager has not been enabled or has been disabled.
 - discovering indicates the subnet manager is sweeping the fabric.
- View the **Activity Count** of the subnet manager. The Activity counter increments each time the subnet manager issues a subnet management packet (SMP) or performs other management activities.
- View the **smKey**. The smkey is a 64-bit subnet management key that is assigned to the subnet manager.
- View the **priority** for the subnet manager. The priority number of a subnet manager tells the subnet manager how to interact with other subnet managers; the highest priority (lowest number) subnet manager becomes the master.
The integer must be between 0 and 15, with the default being 0.
- View the **Sweep Interval** of the subnet manager.
The sweep interval indicates the rate (in seconds) at which the subnet manager sweeps the fabric for any network changes.
The default is 10 seconds.
- View the **Response Timeout** of the subnet manager.
This is the maximum time allowed between the port reception of a subnet management packet and the transmission of the associated response. The default is 2,000 microseconds.

Changing the Subnet Manager Configurations

Change the Priority of a SM

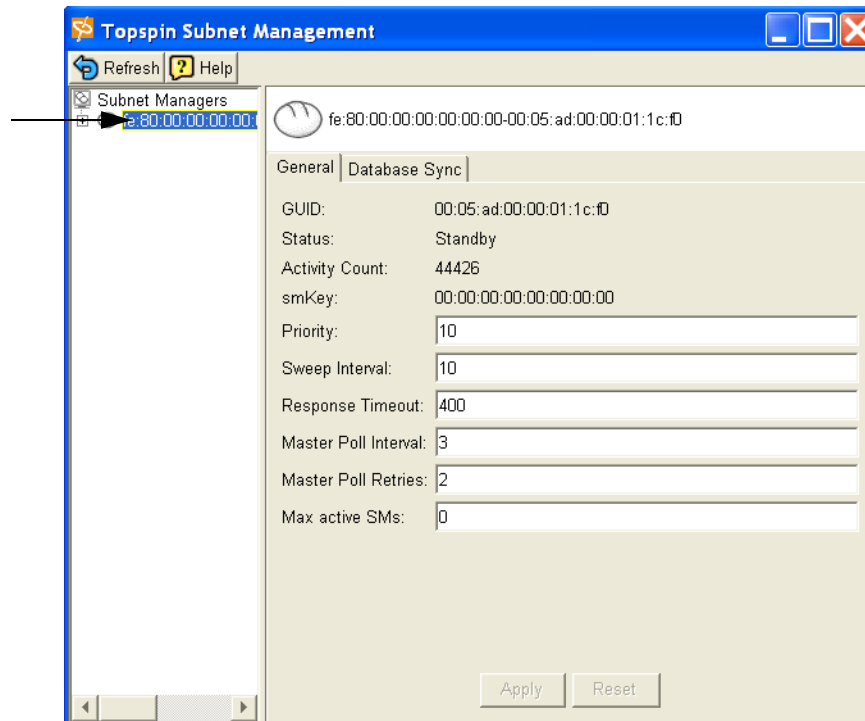
The priority number of a subnet manager tells the subnet manager how to interact with other subnet managers; the highest priority subnet manager becomes the master.

The integer must be between 0 and 15, with the default being 0.

1. Select **InfiniBand > Subnet Management**.

The Subnet Management window appears.

2. Highlight the subnet manager that you want to configure from the left-navigation bar.



3. Click into the **Priority** field.
4. Enter a value between 0 - 15.
5. The **Apply** and **Reset** buttons becomes active if a change is made.
6. Click the **Apply** button to save the changes to the chassis.

Change the Sweep Interval of a SM

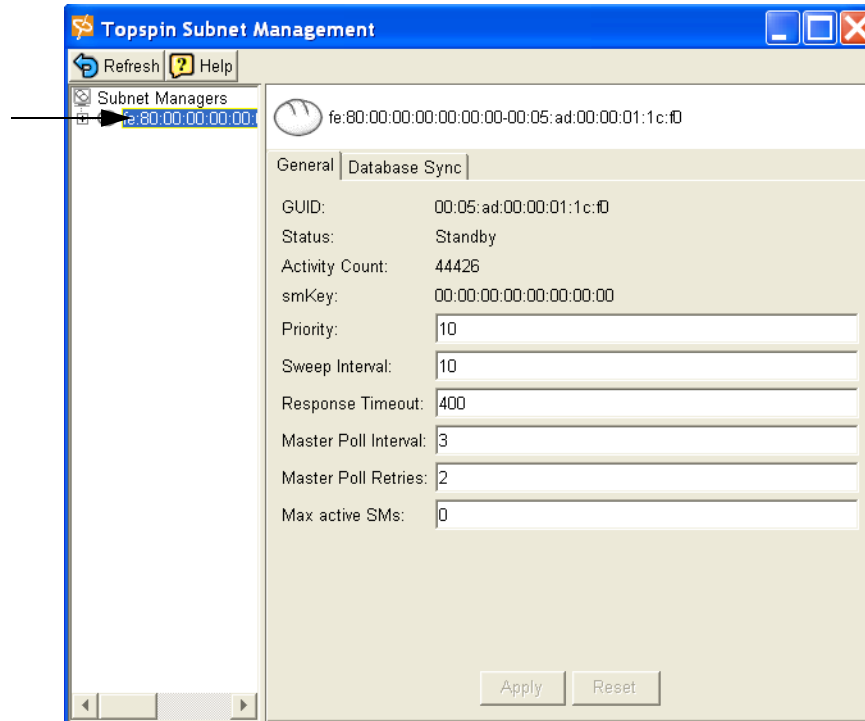
The sweep interval indicates the rate (in seconds) at which the subnet manager sweeps the fabric for any network changes.

The default is 10 seconds.

1. Select **InfiniBand > Subnet Management**.

The Subnet Management window appears.

- Highlight the subnet manager that you want to configure from the left-navigation bar.



- Click into the **Sweep Interval** field.
- Enter a value that indicates the number of seconds between sweeps of the network. The default is 10 seconds. The **Apply** and **Reset** buttons become active if a change is made.
- Click the **Apply** button to save the changes to the chassis.

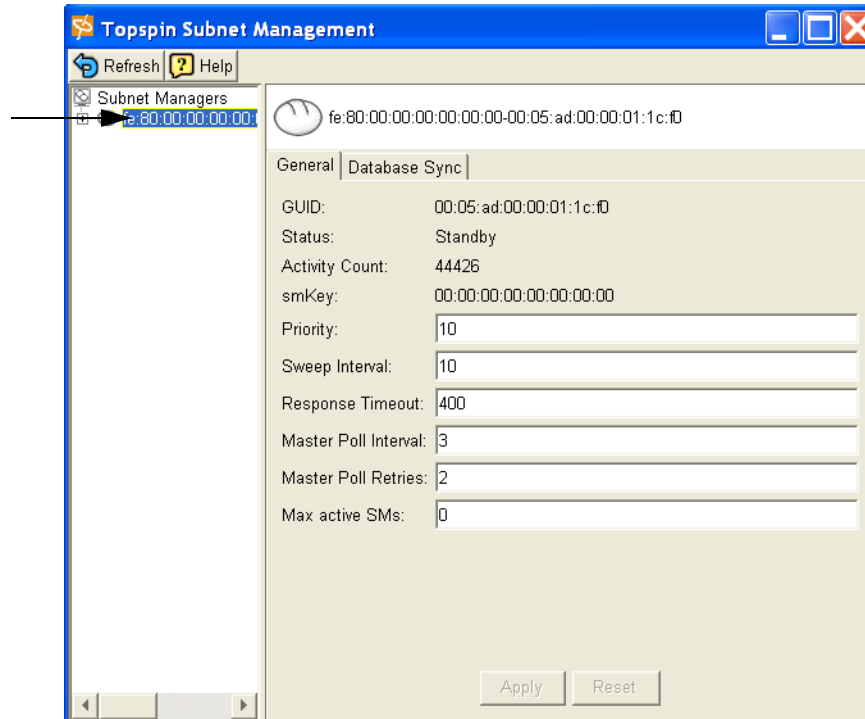
Change the Response Timeout of a SM

The response timeout is the maximum time allowed between the port reception of a subnet management packet and the transmission of the associated response.

The default is 2,000 microseconds.

- Select **InfiniBand > Subnet Management**. The Subnet Management window appears.

- Highlight the subnet manager that you want to configure from the left-navigation bar.



- Click into the **Response Timeout** field.
- Enter a value that indicates the number of microseconds allowed between the port reception of a subnet management packet and the transmission of the associated response.
The **Apply** and **Reset** buttons become active when a change is made.
- Click the **Apply** button to save the changes to the chassis.

Managing Synchronization Between SMs

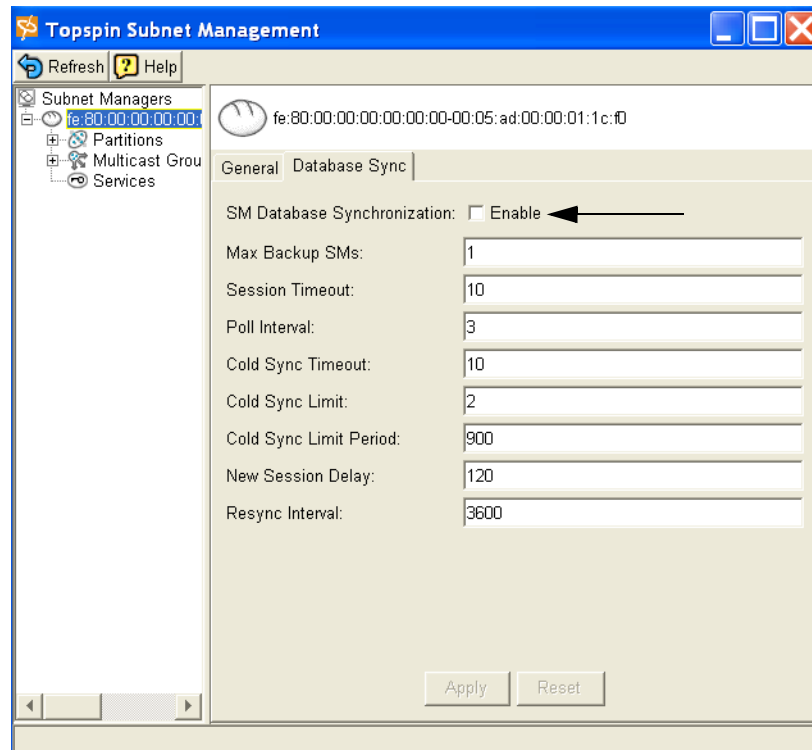
You can enable or disable database synchronization, as well as configure the way database synchronization is performed between the master-Subnet Manager (SM) and one or more standby-SMs. Refer to “[Subnet Manager Hot Standby](#)” on page 5.

Enable/Disable Database Synchronization

Database synchronization is not enabled by default. If you do not enable database synchronization, the contents of the database would be lost whenever a new node assumes the master role in a subnet.

- Select **InfiniBand > Subnet Management**.
The Subnet Management window appears.
- Highlight the subnet manager that you want to configure from the left-navigation bar.
The Subnet Management **General** tab appears.

3. Click the **Database Sync** tab.




4. Click the **Enable** box to enable database synchronization between the active and backup subnet managers.

Set Configurations for the Master SM

1. Select **InfiniBand > Subnet Management**.
The Subnet Management window appears.
2. Highlight the subnet manager that you want to configure from the left-navigation bar.

The Subnet Management **General** tab appears.

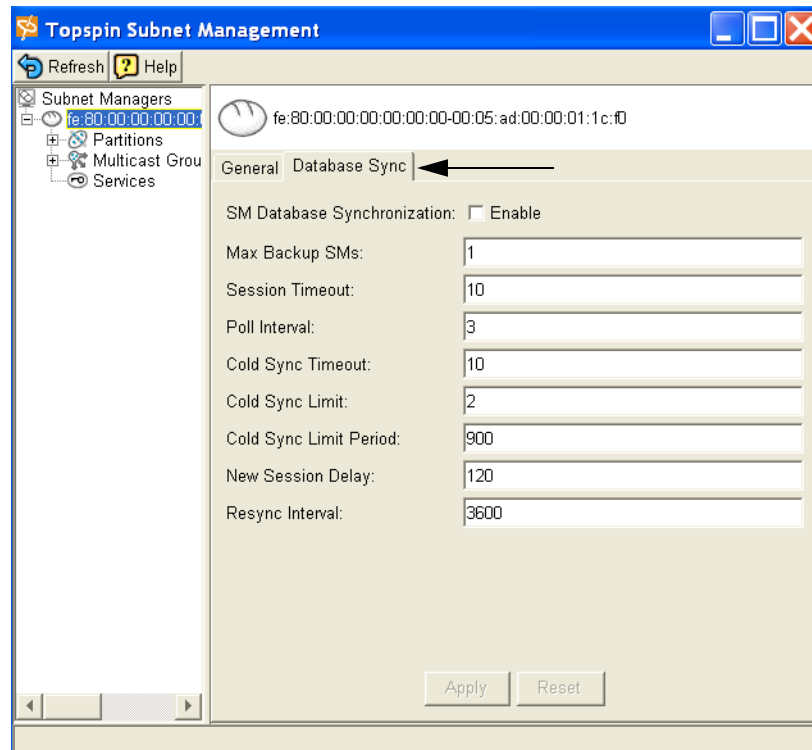


3. Click into the **Master Poll Interval** field to change the interval (in seconds) at which the master SM polls an active slave SM to verify synchronization.
4. Click into the **Master Poll Retries** field to specify the number of unanswered polls that cause the slave to identify the master as dead.
5. Click into the **Max active SMs** field to specify the maximum number of standby SMs that the master supports. Backup SMs are not considered “active.” To set a maximum number of backup SMs, refer to [“Set Configurations for the Backup SM” on page 68](#).
6. Click the **Apply** button to save changes.

Set Configurations for the Backup SM

1. Select **InfiniBand > Subnet Management**.
The Subnet Management window appears.
2. Highlight the subnet manager that you want to configure from the left-navigation bar.
The Subnet Management **General** tab appears.

3. Click the **Database Sync** tab.



4. Click into the **Max Backup SMs** field to enter the maximum number of backup subnet managers with which the master subnet manager will synchronize. A backup subnet manager is automatically added whenever a new InfiniBand (IB) switch is connected to the IB fabric.
The default is 1.
5. Click into the **Session Timeout** field to specify the interval, in seconds, during which a synchronization session status MAD packet must arrive at the master SM to maintain synchronization.
The default is 10 seconds, and the possible entries are 1...30 seconds.
6. Click into the **Poll Interval** field to change the interval at which the master SM polls an active slave SM to verify synchronization.
The default is 3 seconds and the possible entries are 1...30.
7. Click into the **Cold Sync Timeout** field to allot a maximum amount of time in which to perform a cold sync. During the cold sync, the master SM copies all out-of-sync tables to the standby.
The default is 10 seconds and the possible entries are 1...30.
8. Click into the **Cold Sync Limit** field to specify the maximum number of cold syncs that may take place during the cold sync period (see Cold Sync Limit Period).
The default is 2 and the possible entries are 1...10.
9. Click into the **Cold Sync Limit Period** field to specify the length of the interval (in seconds) during which cold syncs may occur. Interval length, in seconds.
The default is 900 seconds and the possible entries are 1...86400 seconds.
10. Click into the **New Session Delay** field to specify the delay (in seconds) before attempting to initiate a synchronization session with a new SM.
The default is 120 seconds and the possible entries are 1...86400 seconds.
11. Click into the **Resync Interval** field to set the interval (in seconds) at which the master will send a re synchronization request to all active synchronization sessions.
The default is 3600 seconds and the possible entries are 1...86400 seconds.

- Click the **Apply** button to save the changes.


Adding a Subnet Manager

Adding additional subnet managers (in addition to the one that is provided by default on the InfiniBand system) should only be done by experienced users.

In the event that additional switch is added to an InfiniBand fabric, an additional subnet manager is added by default (one is the master, and one is the standby).

In most instances, you should use the default subnet manager.

- Open the Subnet Management window by selecting InfiniBand > Subnet Management. Refer to [“Viewing the Subnet Manager Configurations”](#) on page 62.



- Click the **Add** button.
The Add a Subnet Manager dialog box appears.

- Enter a subnet prefix for the subnet manager in the **Subnet Prefix** field.
- Enter a priority number for the subnet manager in the **Priority** field. The value is an integer between 0 (the default) and 15.
The priority number tells the subnet manager how to interact with other subnet managers; the highest priority (lowest number) subnet manager becomes the master.
- Use the default subnet management key in the **smKey** field, which is 00:00:00:00:00:00:00:00:00. The smkey is a 64-bit subnet management key that is assigned to the subnet manager.


- Click the **Add** button.

Viewing Partitions

For complete partition information, refer to [“About Partitions” on page 53](#).


The partitions that are currently being managed by the subnet manager can be viewed by performing the following steps.

- Launch the Element Manager, if you have not already done so.
- Select **InfiniBand > Subnet Management**
The Subnet Management window appears.
- Click open the subnet manager for which you want to view the partitions from the left-navigation tree.
The subnet manager-specific information appears.
- Click on **Partitions** from the left-navigation tree.
The Partition Summary window appears.



- Click on a specific partition from the left-navigation tree.

The Available Members window appears.



About InfiniBand Multicast Groups

An InfiniBand Multicast Group is a collection of Host Channel Adapter (HCA) ports that receive multicast packets sent to a single address.


The configuration and members of a multicast group can be viewed through the Element Manager, but cannot be modified through these screens.

Viewing Multicast Groups

View a Multicast Group Summary

1. Launch Element Manager, if you have not already done so.
2. Select **InfiniBand > Subnet Management**
The Subnet Management window appears.
3. Click open a subnet manager from the left-navigation tree.
The navigation tree opens.
4. Click on **Multicast Groups** from the navigation tree.

The Multicast Groups Summary page appears.



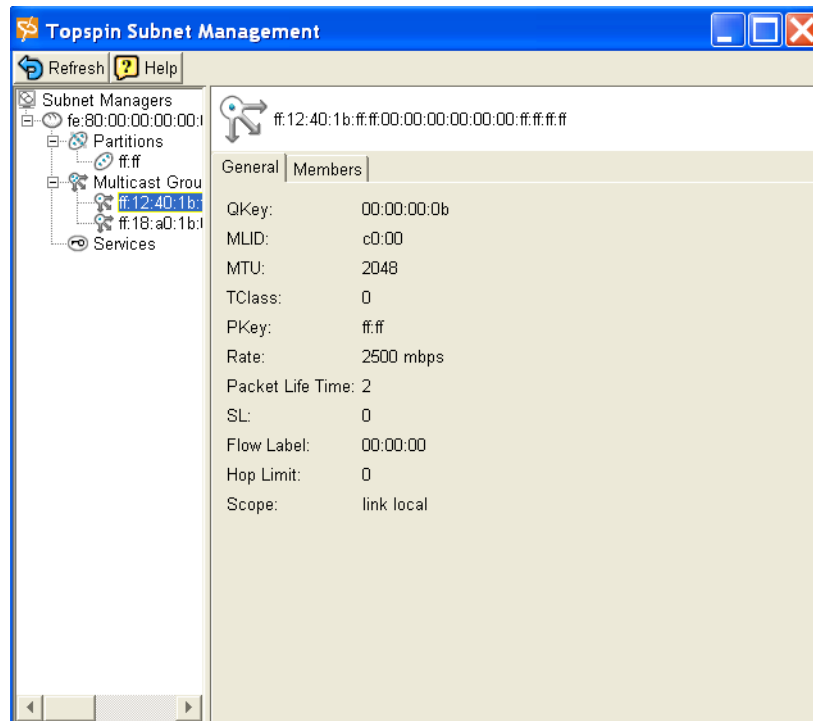
5. View the Multicast Global ID (**MGID**), which is the 64-bit multicast GID address for the multicast group.
6. View the Queue Key (**Q_Key**), which is the 16-bit Q-Key of this multicast group. The queue key is a construct that is used to validate a remote sender's right to access.
7. View the Maximum Transmission Unit (**MTU**) for the multicast group.
8. View the partitions to which the multicast group belongs in the **PKey** field.

View Multicast Group Details

Using the General Tab

1. Launch Element Manager, if you have not already done so.
2. Select **InfiniBand > Subnet Management**
The Subnet Management window appears.
3. Click open a subnet manager from the left-navigation tree.
The navigation tree opens.
4. Click open the **Multicast Groups** folder from the navigation tree.
The Multicast Groups Summary page appears.

- Click on a specific multicast group from the left-navigation tree.




- View the **General** tab, which is displayed by default.
 - View the **Q_Key** for this multicast group. The Queue Key (**QKey**) is a 16-bit construct that is used to validate a remote sender's right to access.
 - View the Local Identifier (**MLID**) for this multicast group. The LID is a 16-bit address that is assigned to a port by the subnet manager. It is used to direct packets within the subnet.
 - View the Maximum Transmission Unit (**MTU**) for the multicast group.
 - View the **TClass** for the multicast group. Specifies the TClass to use in the Global Route Header (GRH), if one is used. A GRH is used in packets that are assigned to destinations outside of a sender's local subnet.
 - View the partitions to which the multicast group belongs in the **PKey** field. Refer to [“About Partitions” on page 53](#).
 - View the traffic **rate** for the multicast group.
 - View the **packet life time** of the multicast group.
 - View the Service Level (**SL**) of the multicast group. The Service Level value is located in the Local Route Header of a packet. It identifies the appropriate virtual lane for a packet, which enables the ability to have multiple services on one physical lane.
 - View the **Flow Label** for the multicast group. This indicates the flow label to be used in the packet's Global Route Header (GRH), if one is used. A GRH is used in packets that are assigned to destinations outside of a sender's local subnet.
 - View the **Hop Limit** for the multicast group. The Hop Limit indicates the limit to be used in the packet's Global Route Header (GRH), if a GRH is used. A GRH is used in packets that are assigned to destinations outside of a sender's local subnet.
 - View the allowable **Scope** of the multicast group.
- Continue to [“Using the Members Tab” on page 74](#).

Using the Members Tab

The Members tab displays the multicast group members and the properties of those members.

- Click the **Members** tab in the Multicast Group window.



- View the **Members** tab:
 - View the Port Global Identifier (**Port GID**) of the multicast group member.
 - View the **Join State** of the multicast group member. The join state may be one or more of the following values: Full Member, Non-Member, and Send Only Member.
 - View the **Proxy Join State** of the multicast group member. The join status can be either True or False.

View the Subnet Manager Services

Services represent actions or functions that a node can perform across the network at the request of another node. Nodes register their services with the subnet manager so other nodes can discover and use these services. The Global Identifier (GID) of a service is the GID of the host that provides the service. Services are mostly used by the DAPL protocol for Address Transferrable Services (ATS), but may also be used by the SRP protocol.


Switch information may be reported for all the switches on a subnet or for a specific switch.

View a Summary of the SM Services

To view the services that are managed by the subnet manager:

- Launch the Element Manager, if you have not already done so.
- Select **InfiniBand > Subnet Management**.
The Subnet Management window appears.
- Click open a subnet manager from the left-navigation tree.

4. Click on **Services** from the left-navigation tree.
The Services window appears.

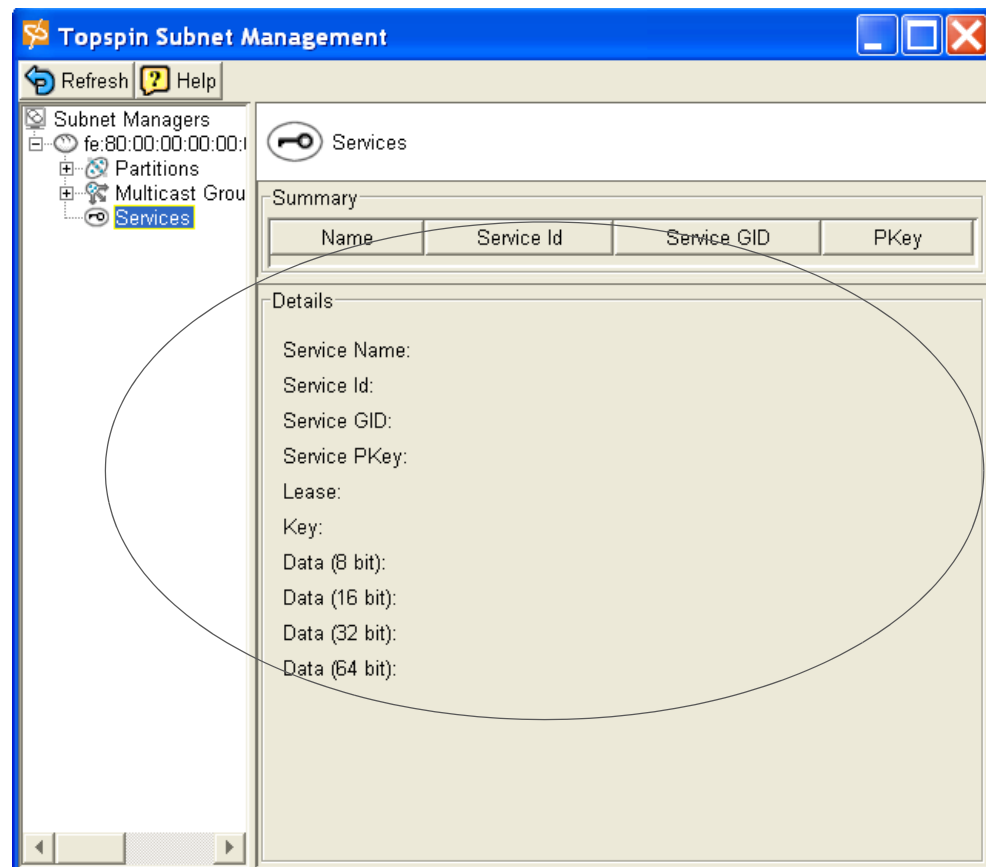


5. View a **Summary** of the selected subnet managers services:
 - View the **Name** of the Service.
 - View the 64-bit ID of the service.
 - View the 128-bit Global ID of the service.
 - View the partition keys affiliated with the service in the **PKey** field. Refer to [“About Partitions”](#) on page 53.

View Details of the SM Services

Details of the subnet managed Services are displayed in the Services window. Refer to [“View the Subnet Manager Services”](#) on page 75.

- View the Details of the subnet managed services:



- View the **Service Name** of the subnet managed service.
- View the 64-bit Service ID of the subnet managed service in the **Service ID** field.
- View the 128-bit Service Global ID of the subnet managed service in the **Service GID** field.
- View the partitions associated with this service in the **Service PKey** field.
- View the lease period remaining (in seconds) for this service in the **Lease** field. The value may be *Indefinite*.
- View the 64-bit service key in the **Key** field.
- View the 8-bit data values associated with this service in the **Data (8-bit)** field.
- View the 16-bit data values associated with this service in the **Data (16-bit)** field.
- View the 32-bit data values associated with this service in the **Data (32-bit)** field.
- View the 64-bit data values associated with this service in the **Data (64-bit)** field.


Configure Subnet Manager Routing

For detailed information regarding Subnet Manager or InfiniBand routing, refer to [“Understanding Subnet Manager Routing”](#) on page 5.

Configure the LID Mask Control (LMC)

The Subnet Manager (SM) allows an administrator to define the LMC (Local Identifier Mask Control) value per subnet. Once the LMC value has been assigned, the SM routes different paths for each LID (an address assigned to a port) that is associated with the same host port.

1. Launch Element Manager, if you have not already done so.
2. Locate the Source LID and Destination LID.
 - a. Select **InfiniBand > SM**.
The Subnet Manager table appears. The default LMC is 0.
3. Click the **Subnet Manager** tab.
4. Use the scroll bar, if necessary, to locate the **LID Mask Control** field.
5. Click into the **LID Mask Control** field and enter the new LMC integer.

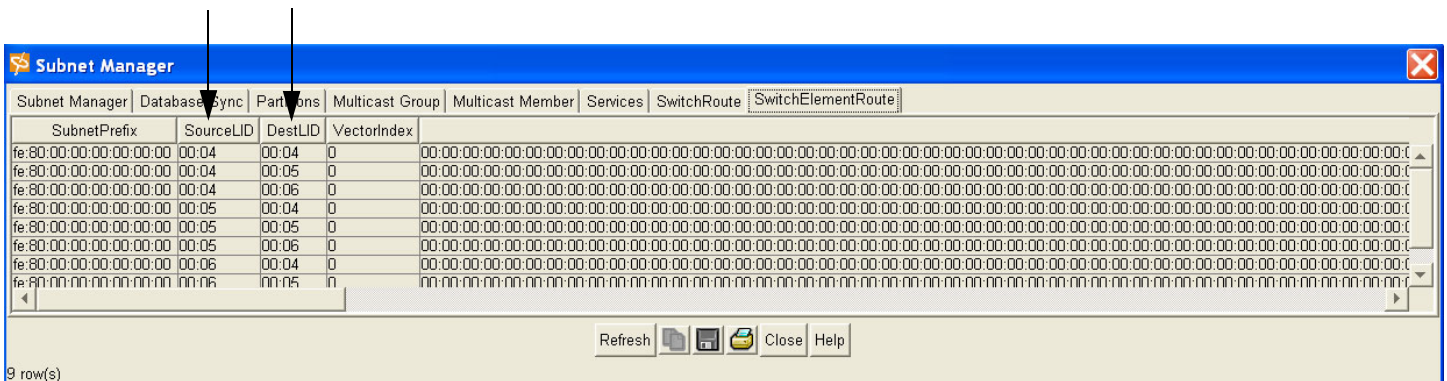


6. Click **Apply**.

View InfiniBand Paths


The following command can be used to help you visualize the path between two end points. The endpoints are specified by Local Identifiers (LIDs):

1. Launch Element Manager, if you have not already done so.
2. Locate the Source LID and Destination LID.
 - a. Select **InfiniBand > SM**.
The Subnet Manager table appears.
 - b. Click on the SwitchElementRoute tab.



3. Return to the Element Manager main menu and select **InfiniBand > Subnet Manager**.
The Subnet Management window appears.
4. Click open the subnet manager from the left navigation tree.
5. Click on Routes from the left navigation tree.

The Routes window appears.



6. Enter the **Source LID** and the **Destination LID** into the fields in the **Filter Route** section.
7. Click the **Show Route** button.

Using the Subnet Manager Through the CLI

This chapter provides the following information:

- [“The Subnet Manager \(SM\)” on page 81](#)
- [“Viewing the Subnet Manager Configurations” on page 82](#)
- [“Changing the Subnet Manager Configurations” on page 83](#)
- [“Managing Synchronization Between SMs” on page 84](#)
- [“Adding a Subnet Manager” on page 87](#)
- [“About InfiniBand Multicast Groups” on page 87](#)
- [“Viewing Multicast Groups” on page 88](#)
- [“Viewing the SM Services” on page 90](#)
- [“Configure Subnet Manager Routing” on page 90](#)

The Subnet Manager (SM)

The subnet manager configures and maintains fabric operations. It is the central repository of all information that is required to setup and bring up the InfiniBand fabric.

Subnet managers are identified by their subnet prefix and Global Unique Identifier (GUID).

There can be multiple Subnet Managers, but only one master.

Master Subnet Manager

The subnet manager that is authoritative, or has the reference configuration information for the subnet.

Standby Subnet Manager

- A subnet manager (SM) that is currently quiescent, and not in the role of a master SM. Standby SMs are dormant managers, and can take over in case of failure of the master subnet manager.

Viewing the Subnet Manager Configurations

View a Summary of Subnet Management

1. Enter the `show ib sm configuration subnet-prefix [prefix | all] summary` command.

Example

```
Topspin-90# show ib sm configuration subnet-prefix fe:80:00:00:00:00:00:00
summary
=====
Subnet Manager Configuration Summary
=====
subnet-prefix          guid                    priority  sm-key
-----
fe:80:00:00:00:00:00:00 00:05:ad:00:00:00:13:f5 10 00:00:00:00:00:00:00:00
Topspin-90#
```

or

Example

```
Topspin-360# show ib sm config subnet-prefix all summary
=====
Subnet Manager Configuration Summary
=====
subnet-prefix          guid                    priority  sm-key
-----
fe:80:00:00:00:00:00:00 00:05:ad:00:00:01:38:82 10          00:00:00:00:00:00:00:00
Topspin-360#
```

An abridged form of the data is displayed. The abridged information includes the subnet prefix, GUID, priority, and SM key of the subnet managers.

View Details of Subnet Management

1. Enter the `show ib sm configuration subnet-prefix [prefix | all]` command.

Example

```
Topspin-90# show ib sm configuration subnet-prefix fe:80:00:00:00:00:00:00
=====
Subnet Manager Information
=====
subnet-prefix : fe:80:00:00:00:00:00:00
guid : 00:05:ad:00:00:00:13:f5
priority : 10
sm-key : 00:00:00:00:00:00:00:00
admin-status : enable
oper-status : master
act-count : 6362
status : active
```

or

Example

```
Topspin-360# show ib sm config subnet-prefix all
```

```
=====
                          Subnet Manager Information
=====
subnet-prefix : fe:80:00:00:00:00:00:00
  guid       : 00:05:ad:00:00:01:38:82
  priority   : 10
  sm-key     : 00:00:00:00:00:00:00:00
admin-status : enable
oper-status  : master
act-count    : 216655
status       : active
sweep-interval : 10
response-timeout : 400
```

- View the subnet-prefix of the subnet manager.
- View the Global Unique Identifier (**guid**) of the subnet manager.
- View the **priority** for the subnet manager. The priority number of a subnet manager tells the subnet manager how to interact with other subnet managers; the highest priority subnet manager becomes the master.
- View the **smKey**. The smkey is a 64-bit subnet management key that is assigned to the subnet manager.
- View the **admin-status** of the subnet manager. The Administrative Status value is enable or disable. Disabling a subnet manager places it in an inactive state but leaves it intact in the configuration. The default is enable.
- View the **oper-status** of the subnet manager. The status is determined by self-detection. The values are notActive, discovering, or Master. As there is only one subnet manager running on the fabric, the sm that is running is always designated the master.
 - notActive indicates the subnet manager has not been enabled or has been disabled.
 - discovering indicates the subnet manager is sweeping the fabric.
- View the **act-count** of the subnet manager. The Activity counter increments each time the subnet manager issues a subnet management packet (SMP) or performs other management activities.
- View the **status** of the subnet manager. The Status may be active or inactive. If active, it is actively managing subnets. If inactive, it is not managing subnets.
- View the **Sweep Interval** of the subnet manager. The sweep interval indicates the rate (in seconds) at which the subnet manager sweeps the fabric for any network changes. The default is 10 seconds.
- View the **Response Timeout** of the subnet manager. This is the maximum time allowed between the port reception of a subnet management packet and the transmission of the associated response. The default is 2,000 microseconds.

Changing the Subnet Manager Configurations

Change the Priority of a SM

The priority number of a subnet manager tells the subnet manager how to interact with other subnet managers; the highest priority subnet manager becomes the master. Because multiple subnet managers

can run on the system and other SMs may run in your IB network, the priority attribute identifies the master SM.

The integer must be between 0 and 15, with the default being 0.

1. Enter the **ib sm subnet-prefix *prefix* priority *sm-priority*** command.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00 priority 10
Topspin-360(config)#
```

Change the Sweep Interval of a SM

The sweep interval indicates the rate (in seconds) at which the subnet manager sweeps the fabric for any network changes.

The default is 10 seconds.

1. Enter the **ib sm subnet-prefix *prefix* sweep-interval *value*** command.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00 sweep-interval
10
```

Change the Response Timeout of a SM

The response timeout is the maximum time allowed between the port reception of a subnet management packet and the transmission of the associated response.

The default is 2,000 microseconds.

1. Enter the **ib sm subnet-prefix *prefix* response time-out *value*** command.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
response-timeout 2000
```

Managing Synchronization Between SMs

You can configure how database synchronization is performed between the master-Subnet Manager (SM) and one or more standby-SMs. Refer to [“Subnet Manager Hot Standby” on page 5](#).

Enable/Disable Database Synchronization

Database synchronization is not enabled by default.

1. Enter the **ib sm db-sync subnet-prefix *<prefix>* enable** command.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm db-sync subnet-prefix fe:80:00:00:00:00:00:00 enable
```


Set Configurations for the Master SM

1. Enter the **config ib sm subnet-prefix <prefix> master-poll-interval <1..60>** command to change the interval (in seconds) at which the master SM polls an active slave SM to verify synchronization. The default is 3 seconds.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
master-poll-interval 1
```

2. Enter the **config ib sm subnet-prefix <prefix> master-poll-retries <1..10>** command to specify the number of unanswered polls that cause the slave to identify the master as dead.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
master-poll-retries 2
```

3. Enter the **config ib sm subnet-prefix <prefix> max-active-sms <0..9999>** command specify the maximum number of standby SMs that the master supports. Backup SMs are not considered “active.”

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
max-active-sms 0
```

Set Configurations for the Backup SM

1. Enter the **ib sm db-sync subnet-prefix <prefix> max-backup-sms <int>** command to enter the maximum number of backup subnet managers with which the master subnet manager will synchronize.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
max-backup-sms 4
```

2. Enter the **ib sm db-sync subnet-prefix <prefix> session-timeout <1..30>** command to specify the timeout in seconds, for receiving a synchronization session status packet, in order to maintain synchronization.

The default is 10, and the possible entries are 1...30 seconds.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
session-timeout 30
```

3. Enter the **ib sm db-sync subnet-prefix <prefix> poll-interval <1..30>** command to change the interval at which the master subnet manager will send a synchronization session status request packet to an active session.

The default is 3 seconds and the possible entries are 1...30.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
poll-interval 1
```

4. Enter the **ib sm db-sync subnet-prefix <prefix> cold-sync-timeout <1..30>** command specify the maximum time in seconds that a cold synchronization should take.

The default is 10 seconds and the possible entries are 1...30.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
cold-sync-timeout 30
```

5. Enter the **ib sm db-sync subnet-prefix <prefix> cold-sync-limit <1..10>** command to allot a maximum amount of time in which to perform a cold sync. During the cold sync, the master SM copies all out-of-sync tables to the standby (see Cold Sync Limit Period).

The default is 2 and the possible entries are 1...10.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
cold-sync-limit 10
```

6. Enter the **ib sm db-sync subnet-prefix <prefix> cold-sync-period** command to specify the maximum number of cold syncs that may take place during the cold sync period (see Cold Sync Limit Period).

The default is 900 and the possible entries are 1...86400.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
cold-sync-period 86400
```

7. Enter the **ib sm db-sync subnet-prefix <prefix> new-session-delay <1..86400>** command to specify the delay (in seconds) before attempting to initiate a synchronization session with a new SM.

The default is 120 and the possible entries are 1...86400.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
new-session-delay 15
```

8. Enter the **ib sm db-sync subnet-prefix <prefix> resync-interval <1..86400>** command to set the interval (in seconds) at which the master will send a re synchronization request to all active synchronization sessions.

The default is 3600 and the possible entries are 1...86400.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:00:00
resync-interval 60
```

Adding a Subnet Manager

Adding additional subnet managers (in addition to the one that is provided by default on the InfiniBand system) should only be done by experienced users.

In the event that an additional switch is added to an InfiniBand fabric, an additional subnet manager is added by default (one is the master, and one is the standby).

In most instances, you should use the default subnet manager.

1. Enter the **ib sm subnet-prefix** *prefix* command.

You must enter a subnet-prefix that does not currently have a subnet manager configured.

Example

```
Topspin-360> enable
Topspin-360# config
Topspin-360(config)# ib sm subnet-prefix fe:80:00:00:00:00:00:01
```

A new subnet manager will be added with the selected subnet-prefix.

2. Configure the attributes for the subnet manager. Refer to [“Changing the Subnet Manager Configurations” on page 83](#).

About InfiniBand Multicast Groups

An InfiniBand Multicast Group is a collection of Host Channel Adapter (HCA) ports that receive multicast packets sent to a single address.

The configuration and members of a multicast group can be viewed through the CLI, but cannot be modified through these screens.

Viewing Multicast Groups

View a Multicast Group Summary

1. Enter the `show ib sm multicast summary` command.

Example

```

Topspin-360> enable
Topspin-360# show ib sm multicast summary
=====
                Summary of Multicast-Groups on Device
=====
subnet-prefix : fe:80:00:00:00:00:00:00
MGID : ff:12:40:1b:ff:f1:00:00:00:00:00:00:ff:ff:ff:ff
multicast-group-members :
port-GID : fe:80:00:00:00:00:00:00:05:ad:00:00:00:12:bf
member-join-state : full-member
proxy-join-status : false
subnet-prefix : fe:80:00:00:00:00:00:00
MGID : ff:12:40:1b:ff:f9:00:00:00:00:00:00:ff:ff:ff:ff

```

- View the multicast group subnet-prefix.
- View the Multicast Global ID (**MGID**), which is the 64-bit multicast GID address for the multicast group.
- View information regarding the multicast group members:
 - View the Port Global Identifier (**Port GID**) of the multicast group member.
 - View the **member-join-state** of the multicast group member. The join state may be one or more of the following values: Full Member, Non-Member, and Send Only Member.
 - View the **proxy-join-status** of the multicast group member. The join status can be either True or False. View the subnet-prefix of the multicast group member.
 - View the Multicast Global ID (**MGID**) of the multicast group member.

View Multicast Group Details

1. Enter the `show ib sm multicast {subnet-prefix prefix | all}` command.

Example

```
Topspin-360# show ib sm multicast subnet-prefix all
```

```
=====
Multicast-Groups Managed by Specific Subnet Manager
=====
subnet-prefix : fe:80:00:00:00:00:00
  MGID : ff:12:40:1b:ff:ff:00:00:00:00:00:00:00:00:00:01
  q-key : 00:00:00:0b
  MLID : c002
  mtu : mtu2048
  t-class : 0
  p_key : ff:ff
  rate : 2500 mbps
packet-life-time : 2
  SL : 0
  flow-label : 00:00:00
  hop-limit : 0
  scope : link-local

  multicast-group-members :
    port-GID : fe:80:00:00:00:00:00:00:05:ad:00:00:00:03:a1
member-join-state : full-member
proxy-join-status : false

    port-GID : fe:80:00:00:00:00:00:00:05:ad:00:00:00:13:18
member-join-state : full-member
proxy-join-status : false

    port-GID : fe:80:00:00:00:00:00:00:05:ad:00:00:00:17:3d
member-join-state : full-member
proxy-join-status : false
```

- View the subnet-prefix of the multicast group.
- View the Multicast Global ID (**MGID**), which is the 64-bit multicast GID address for the multicast group.
- View the Queue Key (**QKey**), which is the 16-bit Q-Key of this multicast group. The queue key is a construct that is used to validate a remote sender's right to access.
- View the Local Identifier (**MLID**) for this multicast group. The LID is a 16-bit address that is assigned to a port by the subnet manager. It is used to direct packets within the subnet.
- View the Maximum Transmission Unit (**MTU**) for the multicast group.
- View the **TClass** for the multicast group. Specifies the TClass to use in the Global Route Header (GRH), if one is used. A GRH is used in packets that are assigned to destinations outside of a sender's local subnet.
- View the partitions to which the multicast group belongs in the **PKey** field. Refer to [“About Partitions” on page 53](#).
- View the traffic **rate** for the multicast group.
- View the **packet life time** of the multicast group.

- View the Service Level (SL) of the multicast group. The Service Level value is located in the Local Route Header of a packet. It identifies the appropriate virtual lane for a packet, which enables the ability to have multiple services on one physical lane.
- View information regarding the multicast group members:
 - View the Port Global Identifier (**Port GID**) of the multicast group member.
 - View the **member-join-state** of the multicast group member. The join state may be one or more of the following values: Full Member, Non-Member, and Send Only Member.
 - View the **proxy-join-status** of the multicast group member. The join status can be either True or False.

Viewing the SM Services

Services represent actions or functions that a node can perform across the network at the request of another node. Nodes register their services with the subnet manager so other nodes can discover and use these services. The Global Identifier (GID) of a service is the GID of the host that provides the service. Services are mostly used by the DAPL protocol for Address Transferrable Services (ATS), but may also be used by the SRP protocol.

Switch information may be reported for all the switches on a specific subnet or for a specific switch.

View a Summary of the SM Services

To view the services that are managed by the subnet manager:

1. Enter the **show ib sm service summary** command.

Example

```
Topspin-360# show ib sm service summary
=====
                          Summary of Services on Device
=====
subnet-prefix : fe:80:00:00:00:00:00:00
  service-id  : 10:00:0c:e1:00:41:54:53
    GID       : fe:80:00:00:00:00:00:00:00:05:ad:00:00:01:29:8d
  service-data :
    data-8    : 00:00:00:00:00:00:00:00:00:00:00:00:0a:04:c4:73
    data-16   : 0000:0000:0000:0000:0000:0000:0000:0000
    data-32   : 00000000:00000000:00000000:00000000
    data-64   : 0000000000000000:0000000000000000
subnet-prefix : fe:80:00:00:00:00:00:00
  service-id  : 10:00:0c:e1:00:41:54:53
    GID       : fe:80:00:00:00:00:00:00:00:05:ad:00:00:01:29:c6
  service-data :
    data-8    : 00:00:00:00:00:00:00:00:00:00:00:00:0a:04:c4:72
    data-16   : 0000:0000:0000:0000:0000:0000:0000:0000
    data-32   : 00000000:00000000:00000000:00000000
    data-64   : 0000000000000000:0000000000000000
```

Configure Subnet Manager Routing

For detailed information regarding Subnet Manager or InfiniBand routing, refer to [“Understanding Subnet Manager Routing”](#) on page 5.

Configure the LID Mask Control (LMC)

The Subnet Manager (SM) allows an administrator to define the LMC (Local Identifier Mask Control) value per subnet. Once the LMC value has been assigned, the SM routes different paths for each LID (an address assigned to a port) that is associated with the same host port.

The default LMC is 0.

View InfiniBand Paths

The following command can be used to help you visualize the path between two end points. The endpoints are specified by Local Identifiers (LIDs):

1. Locate the Source LID and Destination LID.
2. Enter **show ib sm switch-route subnet-prefix** fe:80:00:00:00:00:00:00 **src-lid** <source-lid> **dst-lid** <destination-lid>.

The following example displays a summary of the SM route switch element table for one source and destination.

Example

```
Topspin-90> show ib sm switch-route subnet-prefix
fe:80:00:00:00:00:00:00 src-lid dst-lid 9

=====
Summary of SM Route Switch Element Table by Subnet w/ Src and Dest LID
=====
subnet-prefix : fe:80:00:00:00:00:00:00
src-lid : 1
dst-lid : 1
last-change : Tue Jan 27 22:51:56 2004
```


Using Image Files

This chapter describes the following Access and Security features:

- [“Types of Image Upgrades” on page 93](#)
- [“Upgrade Procedure Overview” on page 95](#)
- Upgrade Procedure Steps (begins [page 95](#))

Types of Image Upgrades

There are two different types of upgrades that may need to be downloaded to the Topspin system.

TopspinOS Upgrades

Downloads can be performed remotely using an tftp or an ftp server. Once the file has been downloaded, it can be pushed to the system through the Element Manager or through the CLI.

The following chapter describes the upgrade procedure for the TopspinOS (refer to [page 95](#)).

About the System Image

The Image data that is used to configure the software is being continuously updated and enhanced. Use the latest system image data to ensure the most efficient usage of your system.

What is a System Image?

A system-image is an unpacked and installed image file. An image file is the source from which to install a system-image and it has a `.img` extension.

When an image file is installed, the image file is expanded into a “system image.” The system image is what the user will refer to in order to specify what the system should use to boot-up each card in the system.

What is an Image File?

Image files are stored in flash memory as a single complete file with a “.img” extension. Each image file contains all the operating software (application software and firmware/microcode) needed by the various cards that can be installed into the Topspin system.

The Topspin system cannot use an image file directly to boot-up to system. The image file must first be installed. The installation process automatically unbundles the image file and distributes the software components to each card in the system. Users do not have to be aware of individual software components. The user simply executes one CLI command to install an image file. Refer to the **install** command in the *HP 24-Port 4x Fabric Copper Switch Command Line Interface Reference Guide*.

The TopspinOS stores up to three images on a disk: the uninstalled image, the current system (or “installed”) image, and the recovery image.

The system only has enough flash memory to store:

- one system image file (active)
- one image file (inactive/uninstalled)
- one recovery image

Occasionally, you will have to manually delete an image file from the InfiniBand system to make room for a new version. Refer to [“Deleting Image Files” on page 103](#).

Inactive Image

An inactive image is one that has been downloaded, but has not been installed; therefore it is not the active, or “system” image.

The TopspinOS can only store one inactive image. Delete inactive images through the CLI (refer to [“Deleting Image Files” on page 103](#)), or by clicking the **delete** button in the Element Manager.

Active Image

The term “Active Image” refers to the current System Image. An installed, or “active” image has gone through the entire upgrade process. The System-Image usually has a slash (/) in its name. Do not modify or delete the installed system-image.

Recovery Image

The Recovery Image is a default image that comes installed on the Topspin system. The Recovery Image can be used to quickly restore operation to the system if an image upgrade should fail.

About Copying/Downloading the Image

Upgrading the TopspinOS requires several steps, which are described in the following sections. One of the steps will be to copy the image before installing it.

There are several options when copying the image into the system:

Table 9-1: Copying/Downloading Image Options

Through the CLI	Through the GUI
FTP	Remote FTP Server
TFTP	Local File

Note: Images cannot be installed through the GUI. Once the image has been copied via the chosen method, refer to [“Activate an Image” on page 100](#).

Card Status Requirements

Only cards with an oper-status of “up” are updated. If a card is down when you run **install**, or a card is added after running **install**:

1. Bring up the card
2. Run **install** again.
Specify the same image file. If the image is already installed on a card, that card is skipped.
3. Be sure to specify the **boot-config** again so that all cards know to boot from the same system image.

Upgrade Procedure Overview

The system upgrade process is summarized in the following steps:

1. Set up the hardware connection for the upgrade ([page 95](#)).
2. Verify the installed system-image version number ([page 96](#)).
3. Download an image file ([page 96](#)):
 - from a network-accessible ftp server.
 - or
 - Download an image file remotely from a tftp server.
4. Install the new system-image ([page 100](#)).
5. Configure the CLI and Element Manager to use the appropriate configuration file the next time they reboot ([page 101](#)).
6. Reboot the system ([page 102](#)).

Set-Up the Hardware Connection

For detailed information about the hardware, refer to the *HP Dual-port 4x Fabric Adapter User Guide* and the *HP 24-Port 4x Fabric Copper Switch Hardware User Guide*.

There are two types of hardware connections that can be used to download a new image to the InfiniBand system: Out-of-Band or In-Band.

Out-of-Band Connection

1. Connect the server to the ethernet management port, if a connection does not already exist.


In-Band Connection

1. Use a 4x copper cable to connect from an InfiniBand-enabled host to an InfiniBand port on the InfiniBand switch.

Verify the Installed Image Version

Check the Image Version Through the GUI

1. In the Element Manager, select **Maintenance** -> **File Management...**
The File Management Window appears.



2. Note the version number of the installed image by looking in the **File Name** column.
The installed system image has a **.img** extension.
3. Image files that are *not* installed files (for example, the **.img** files) can be deleted at this time to make room for the latest version by using the **delete** command.
4. Highlight the name of an uninstalled image file and click the **Delete** button.
Note: Do not delete an installed image file, which will have a slash (/) in the name.

Check the Image Version Through the CLI

1. Enter the **dir image** command to show all the current images on the system.

```

Topspin-360> enable
Topspin-360# dir image
=====
                        Existing Boot-Images on System
=====
slot date-created          size      file-name
-----
1    Tue Nov 25 18:46:32 2003    30695264    Topspin360-TopspinOS-2.0.0-build212.img
1    Tue Nov 25 19:34:08 2003    19228160    TopspinOS-2.0.0/build211
1    Tue Nov 25 19:32:16 2003    15539200    TopspinOS-2.0.0/build212
Topspin-360#

```

2. Note the version number of the installed image by looking for the file that has a **.img** extension.
3. Image files that are *not* installed files (for example, the **.img** files) can be deleted at this time to make room for the latest version by using the **delete** command.
Note: Do not delete an installed image file, which will have a slash (/) in the name.

Copy/Download the Image

- [“Copy/Download the Image Through the GUI” on page 97](#)
or

- [“Copy/Download an Image Through the CLI” on page 98](#)

Copy/Download the Image Through the GUI

Images must be installed using the CLI; however, images can be *copied* from a remote or local location through the GUI.

Copy an Image from a Remote Location

This section describes using the Element Manager to copy an image file from a remote location before installing it as the active image.

1. In the Element Manager, select **Maintenance** -> **File Management...**

Note the uninstalled image files(.img) and installed system-images. Do not modify or delete the *installed* system-image. If there are two saved image files, you can create space by clicking **Delete**.

2. Click the **Import...** button.

The Import File window appears.

3. Select *Image* from the **File Type** drop-down menu.
4. Click the **Remote FTP** radio button from the **Copy From** section.
5. Enter the IP address of the FTP server from which to copy the file in the **Server Name or IP Address** field.
6. Enter the network name of a recognized user in the **User Name** field.
7. Enter the password for the specified user in the **Password** field.
8. Enter the path and name of the image file on the FTP server in the **File Path and Name** field. Image file names must include the “.img” extension. The “.cfg” extension is optional when specifying configuration file names.
9. Enter the name of the host to which you want to copy the file in the **Copy To** field.
10. Click the **Copy** button to copy the file or the **Cancel** button to close the window.

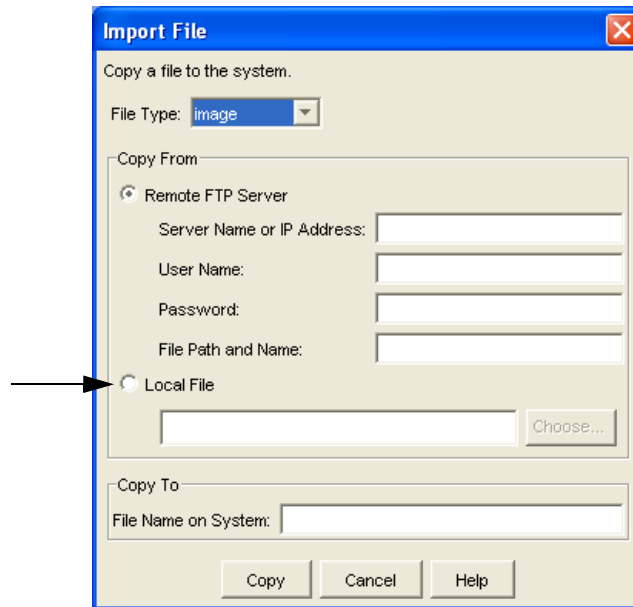
Wait until the transfer is complete.

The window will automatically refresh to show the latest copied image file. However, the .img file will not change until you have installed the new image. Installation must be performed through the CLI.

Copy an Image from a Local Location

This section describes using the Element Manager to copy an image file from a locally available location before installing it as the active image.

1. In the Element Manager, select **Maintenance** -> **File Management...**
Note the uninstalled image files(.img) and installed system-images. Do not modify or delete the *installed* system-image. If there are two saved image files, you can create space by clicking **Delete**.
2. Click the **Import...** button.
3. Copy/Download an Image Through the CLI
The Import File window appears.



4. Select *Image* from the **File Type** drop-down menu.
5. Click the **Local File** radio button from the **Copy From** section.
6. Click the **Choose...** button.
7. Navigate to the locally available image file and select the file.
8. Enter the file name of the image in the **Copy To: File Name on System** field.
Image files must reside in the image file-system and the file name must have the *.img* extension.
9. Click the **Copy** button.
Wait until the transfer is complete.
The window will automatically refresh to show the latest copied image file. However, the *.img* file will not change until you have installed the new image. Installation must be performed through the CLI.

Copy/Download an Image Through the CLI

There are two ways that the software can be downloaded from the CLI:

- Through an FTP server
- Through a TFTP server

Download From an FTP Server

Use FTP to download new image files to the InfiniBand system. Uninstalled system-image files always end with a *.img* extension.

Directory management is automatically performed on the InfiniBand system, so do not include path information for files on the chassis.

1. (Optional) If you are using In-Band Management, configure the ethernet management interface (if you have not already done so).

```
Topspin-90> enable
Topspin-90# config
Topspin-90 (config)# interface mgmt-ether
Topspin-90 (config-if-mgmt-ethernet)#ip address 10.3.102.8. 255.255.0.0
Topspin-90 (config-if-mgmt-ethernet)#gateway 10.3.0.1
Topspin-90 (config-if-mgmt-ethernet)#no shutdown
Topspin-90 (config-if-mgmt-ethernet)# exit
Topspin-90 (config) # exit
```

2. Verify that you have a working ftp server connection, and that you have a user account on the ftp server.

```
$ftp 10.10.0.5
username:Example
password:xxxxxx
ftp>
```

3. Enter the **copy** command in the privileged-execute mode with the source, destination, and FTP user information.

Syntax:

```
Topspin-90> enable
Topspin-90# copy
ftp://user-name:password@source-IP-address/source-image-file-path-name
image:destination-file-name
```

Enter path, user-name and password information.

This information authenticates you to the server, which is specified in source-IP-address.

- Include the full path to the file.
- Set the file-system to image when copying image files.

Example:

```
Topspin-90# copy
ftp://bob:mypassword@10.0.0.5/Topspin-360-TopspinOS-1.1.3-build497.img
image:Topspin-360-TopspinOS-1.1.3-build497.img
*****operation completed successfully
Topspin-90#
```

4. Continue on to [“Activate an Image” on page 100](#) to install the image.

Download From a TFTP Server

Use FTP to download new image files to the InfiniBand system. Uninstalled system-image files always end with a .img extension.

Directory management is automatically performed on the InfiniBand system, so do not include path information for files on the chassis.

1. (Optional) If you are using In-Band Management, configure the ethernet management interface (if you have not already done so).

```
Topspin-270> enable
Topspin-270# config
Topspin-270 (config)# interface mgmt-ether
Topspin-270 (config-if-mgmt-ethernet)#ip address 10.3.102.8. 255.255.0.0
Topspin-270 (config-if-mgmt-ethernet)#gateway 10.3.0.1
Topspin-270 (config-if-mgmt-ethernet)#no shutdown
Topspin-270 (config-if-mgmt-ethernet)# exit
Topspin-270 (config) # exit
```

2. In the privileged-execute mode, enter the **copy** command with the remote system and path information.

The “remote system” must be the IP address of the TFTP server.

Syntax:

```
Topspin-270> enable
Topspin-270# copy tftp://<remote-system>/<filepath> <filesystem>:<filename>
```

Example:

```
Topspin-270# copy tftp://@10.10.20.78/tftp_dir/Topspin270-TopspinOS.img
image:Topspin270-TopspinOS-2.0-build397.img
*****operation completed successfully
Topspin-270#
```

3. (Optional) View the images by using the **dir image** command.

```
Topspin-270# dir image
=====
Existing Boot-Images on System
=====
slot date-created          size      file-name
-----
1   Mon Aug 11 22:41:19 2003  29830309
Topspin270-TopspinOS-2.0-build397.img
```

4. Continue on to [“Activate an Image” on page 100](#) to install the image.

Activate an Image

After downloading the image file to the chassis controller, it must be installed to become active. The **install** command installs the specified image file into the system.

To install an image file:

1. In the privileged-execute mode, enter **install image:name of file**
 - Image files must reside in the image file-system
 - The file name must have the *.img* extension
 - All cards must have an oper-status of “Up”

Syntax:

```
Topspin-360#install image:image-file-name.img
```



```

Topspin-360# install image:Topspin-360-TopspinOS-1.1.3-build497.img
*****operation completed successfully
Topspin-360#

```

The name of the image file will change after it is installed.

2. Repeat the install procedure on any cards that did not have an oper-status of “up” at the time of image installation.

View a card’s status by using the **show card** command.

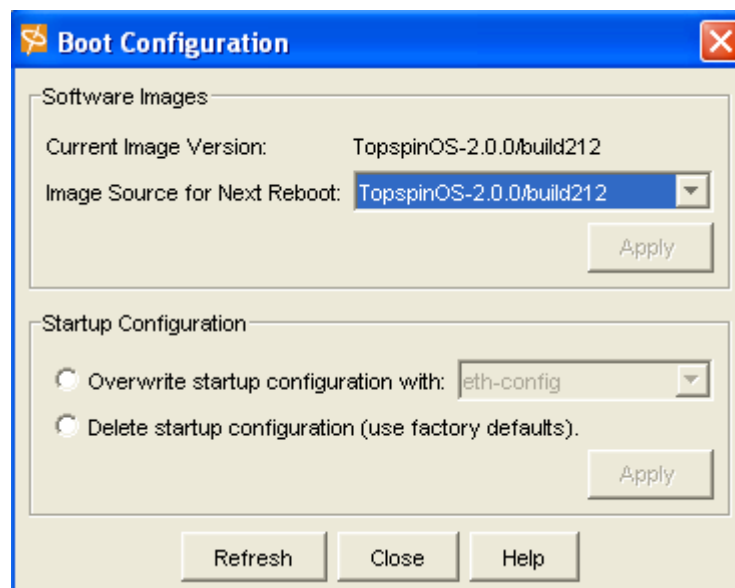
Specify a New Boot Image

After a system-image is installed on the system chassis, specify the system-image to use the next time it reboots.

- [“Specify a New Boot Image Through the GUI” on page 101](#)
- [“Specify a New Boot Image Through the CLI” on page 102](#)

Specify a New Boot Image Through the GUI

1. In the Element Manager, select **Maintenance-> Boot Config...**
The Boot Configuration window appears.



2. Select the new image configuration from the **Image Source for Next Reboot** drop-down menu.
3. Click the **Apply** button.

Specify a New Boot Image Through the CLI

1. Find the name of your new system image by listing all the image files on the Topspin chassis. Enter the **dir image** command.

```
Topspin-360# dir image
=====
Existing Boot-Images on System
=====
slot date-created          size      file-name
-----
1    Tue Feb 25 16:14:15 2003  23198989  TopspinOS-2.0-build491
1    Thu Jan  1 00:18:52 1970   1024      TopspinOS-2.0/build491
1    Tue Feb 25 16:17:04 2003   1024      TopspinOS-2.0/build497
Topspin-360#
```

2. Copy the displayed system-image (this will be pasted into the **boot-config** command).
3. Specify the new boot image.
Enter **boot-config primary-image-source**, then paste the then paste the copied *image* file name.

Syntax:

```
boot-config primary-image-source image-directory-name
```

Example

```
Topspin-360# configure
Topspin-360(config)# boot-config primary-image-source
TopspinOS-1.1.3/build497
Topspin-360(config)# exit
```

4. Show the new system -image.

Example

```
Topspin-360# show boot-config
=====
System Boot Configuration
=====
slot-id : 1
sw-version : TopspinOS-1.1.3/build491
last-image-source : TopspinOS-1.1.3/build491
primary-image-source : TopspinOS-1.1.3/build497
```

Reboot the System

The complete reboot process is described in the other chapters.

- Reboot the system through the GUI (refer to [page 37](#)).
- Reboot the system through the CLI (refer to [page 43](#)).

Deleting Image Files

Up to two image files may be saved on the system. Older versions of image files will have to be removed before you can add more.

- [“Deleting Images Through the GUI” on page 103](#)
- [“Deleting Images Through the CLI” on page 103](#)

Deleting Images Through the GUI

For information regarding deleting images through the Chassis Manager web GUI, refer to the *HP 24-Port 4x Fabric Copper Switch Chassis Manager User Guide*.

1. In the Element Manager, select **Maintenance** -> **File Management....**
The File Management window appears.
Note the uninstalled image files(.img) and installed system-images.
2. Click on the name of the extraneous uninstalled image that you wish you delete in the **File Name** field.
Note: Do not modify or delete the installed system-image.
3. Click the **Delete** button.
A prompt appears to verify that you want to delete the selected image.
4. Click the **Yes** button.

Deleting Images Through the CLI

1. Enter the **delete image** command and the name of the extraneous uninstalled image that you wish you delete in the Privileged Execute mode.
Enter the file name exactly as it is displayed by the **dir** command.

Syntax:

```
delete image:file
```

Example

```
Topspin-90# delete image:Topspin90-TopspinOS-2.0.0-build211.img
*****
Topspin-90#
```

You will be prompted to confirm that you want to delete the file.

Example

```
Delete file TopspinOS-2.0.0-build211.img? [yes(default) | no] Y
Topspin-90#
```

2. Enter *Y* to delete the file.
3. Press the Enter key.

Using Configuration Files

This chapter describes:

- [“Understanding Configuration Files” on page 105](#)
- [“Listing Configuration Files” on page 106](#)
- [“Export a Configuration File” on page 106](#)
- [“Import a Configuration File” on page 108](#)

Understanding Configuration Files

A configuration file is a text file that stores a list of CLI commands.

About the Startup-Config

The main configuration file is called `startup-config`. This file stores all of the CLI commands necessary to completely configure a box from a factory, default state. This configuration file can be copied, backed up, and modified.

About the Running-Config

Whenever configuration changes are made via the GUI or CLI, a CLI command is temporarily saved in a virtual configuration file called `running-config`. If the administrator wishes to save these changes permanently, this file is “copied” into the `startup-config` file.

Any number of configuration files can be stored. For convenience and rapid configuration, files can also maintain a partial list of CLI commands. These can also be copied into `running-config` for immediate use or `startup-config` for persistent use across reboots.

Listing Configuration Files

List Config Files Through the CLI

To list the configuration files currently stored on the InfiniBand system, enter the `dir` command with the `config` keyword.

For example:

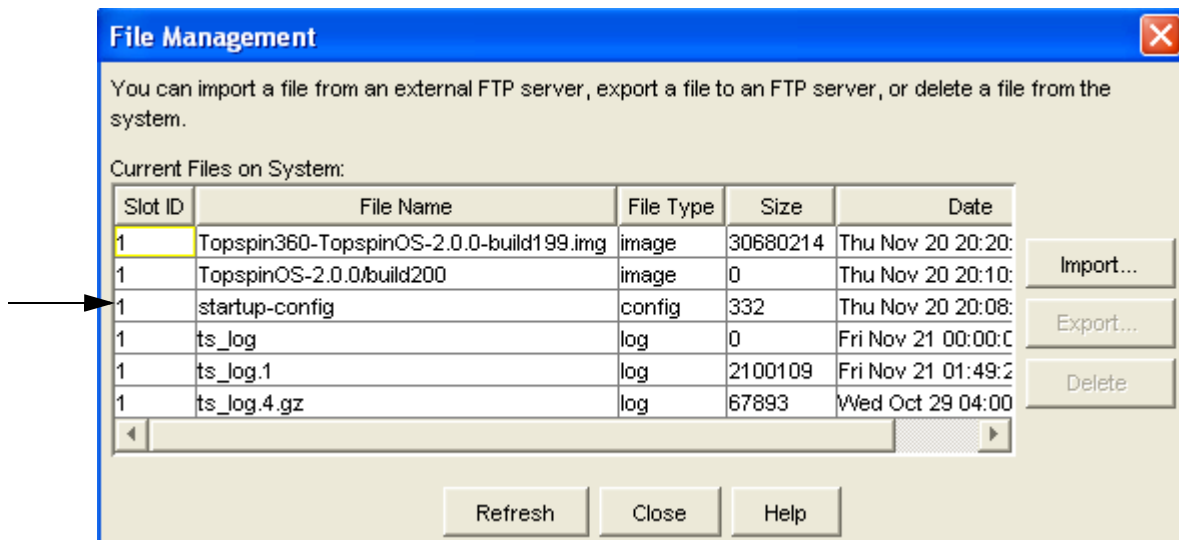
```
Topspin-360# dir config
=====
Existing Configurations on System
=====
slot date-created          size  file-name
-----
1    Tue Jan 14 23:19:55 2003  9110   startup-config2
1    Tue Jan 14 00:11:04 2003 13925   startup-config
Topspin-360#
```

List Config Files Through the GUI

To show a list of current configuration files through the Element Manager:

1. Launch the Element Manager, if you have not already done so.
2. Select **Maintenance --> File Management**.

The File Management window opens.



3. View the **File Name** or the **File Type** columns. The configuration files are called "config."

Export a Configuration File

Upload configuration and log files to maintain backups and troubleshoot your device.

Export a Config File Through the CLI

To copy a configuration file to a remote FTP server through the CLI:

1. Enter the **copy** command, as well as the necessary ftp information in one string:
 - ftp file system
 - name of the config file that you want to copy.
 - **ftp** keyword
 - your ftp username and password
 - FTP server domain name or IP address.
 - directory path on the host to which you want to copy the config file.
 - name of the file where you want to store the copied config file.

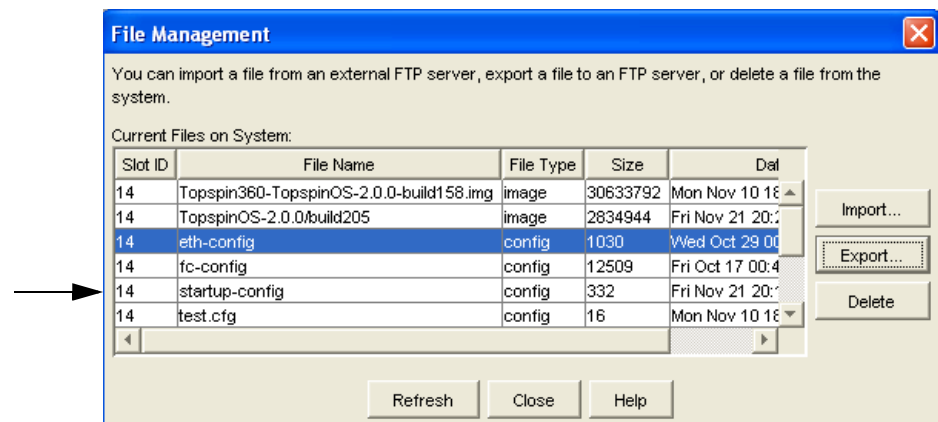
```
Topspin-360# copy sampleserver startup-config.cfg ftp://bob-secret@10.10.2.40
/ftpserver/startup-config.cfg
```

Export a Config File Through the GUI

To copy a configuration file through the Element Manager:

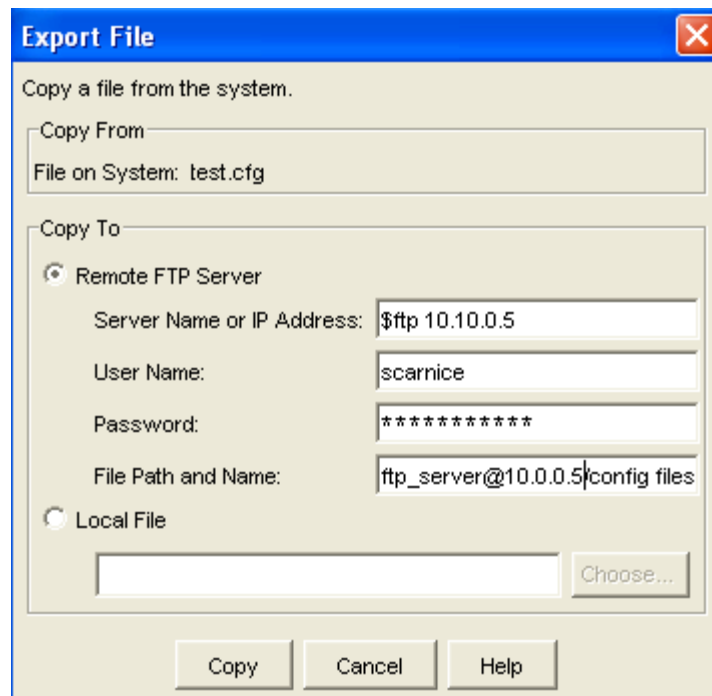
1. Launch the Element Manager, if you have not already done so.
2. Select **Maintenance --> File Management**.

The File Management window opens.



3. Click on the config file that you want to save in the **Current Files on the System** window.
4. Click the **Export** button.
The Import File window opens.
5. Select the method by which you want to export the file:
 - Copy to a remote ftp server.
 - Copy to a local file.

6. (Optional) If you are copying the config file to an ftp server:
 - a. Verify that you have a connection and privileges to an FTP server.
 - b. Click the **Remote FTP Server** radio button in the Copy To section.
 - c. Enter the name or IP address of the FTP server in the **Server Name or IP Address** field.
 - d. Enter your User Name for the FTP server in the **User Name** field.
 - e. Enter your FTP server password in the **Password** field.
 - f. Enter the path to the FTP server, as well as the file name into the **File Path and Name** field.



7. (Optional) If you are copying the config file to a local drive:
 - a. Click the **Local File** radio button in the Copy To section.
 - b. Click the **Choose** button.
The Select File window appears.
 - c. Navigate to the location where you want to store the config file.
 - d. Change the name of the config file in the File Name field, if desired.
8. Click the **Copy** button.

Import a Configuration File

Download configuration files from either a FTP or TFTP server to quickly replicate a desired configuration.

Download a Config File Through the CLI

Download a Config File From a FTP Server

To download a configuration file from a remote FTP server through the CLI:

1. Enter the **copy** command, as well as the necessary ftp information in one string:
 - **ftp** keyword
 - your ftp *username:password*

- @ FTP server domain name or IP address.
- directory path on the host from which you want to copy the config file.
- name of the config file that you want to copy.
- directory path on the host to which you want to copy the config file.
- name of the file where you want to store the copied config file.

```
Topspin-360# copy ftp://bob:secret@10.0.0.5/random directory path/startup.cfg  
sampledirectory:startup.cfg
```

Download a Config File From a TFTP Server

To download a configuration file from a TFTP server through the CLI:

1. Enter the **copy** command, as well as the necessary ftp information in one string:
 - **ftp** keyword
 - your ftp username and password
 - name of the config file that you want to copy.
 - @ FTP server domain name or IP address.
 - path on the host to which you want to copy the config file.
 - name of the config file that you want to import.

```
Topspin-360# copy ftp://bob:samplecompany@10.0.0.5/Topspin-360-Basic.cfg  
image:IB.cfg
```


Using Log Files

This chapter describes:

- [“Understanding Log Files” on page 111](#)
- [“Listing Current Log File Names” on page 112](#)
- [“Viewing a Log File Through the CLI” on page 113](#)
- [“Viewing a Log File Through the GUP” on page 114](#)
- [“Configuring Remote Logging” on page 117](#)

Understanding Log Files

Log files are text files that record activity, including configuration changes. Depending on size, log files are rotated and compressed.

Log files can also be exported from the Topspin system by using the **copy** command.

File Management and Storage

The management of log files is performed automatically, but can be configured. Log files are stored separately from other file types, but all files share the 128 MB of flash memory.

Log files are stored in syslog files.

The system checks the size of the active log file hourly, and when it exceeds 1 MByte, the active log file, `ts_log`, is closed, compressed, and renamed `ts_log.1.gz`. Other `ts_log.x.gz` files are incremented by 1. These files can be downloaded via the Log Viewer GUI, which can create filters for troubleshooting and auditing purposes.

About Message Types

The following levels of logging are captured:

- **CONF** - configuration changes. No user action is required
- **INFO** - general information. No user action is required
- **WARN** - abnormal condition. User intervention may be required
- **ERROR** - abnormal condition. User intervention is required
- **FATAL** - abnormal condition. User must reboot
- **TRACE** - Refer to [“About Tracing” on page 149](#) for information regarding Traces.

Listing Current Log File Names

Listing Current Logs Through the CLI

To list the log files currently stored on the InfiniBand system:

1. Enter the `dir syslog` command.

For example:

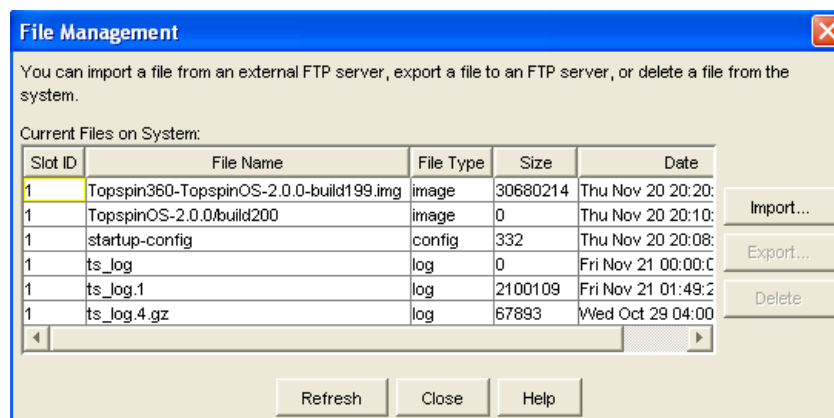
```
Topspin-90# dir syslog
=====
Existing Syslog-files on System
=====
slot date-created          size    file-name
-----
1   Thu Nov 20 20:26:28 2003  204331  ts_log
1   Tue Nov 18 10:00:04 2003   66965  ts_log.3.gz
1   Wed Oct 29 04:00:03 2003   67893  ts_log.4.gz
Topspin-90#
```

Listing Current Logs Through the GUI

1. Launch Element Manager, if you have not already done so.

Select **Maintenance --> File Management**.

The File Management window appears.



2. Look for the log in the **File Type** column.
3. Note the corresponding name of the log file in the **File Name** column.
4. Click in a row that contains a log file.

From this window you can **Export** or **Delete** a log file.

Viewing a Log File Through the CLI

Display Entire Log

1. Set terminal parameters, if you have not already done so. Refer to [“Setting Terminal Parameters” on page 22](#).
2. Enter the **show logging** command.

```

Topspin-360# show logging
Nov 19 22:08:49 topspin-cc kernel: THH kernel module initialized successfully
Nov 19 22:09:17 topspin-cc ts_sma.x[746]: [INFO]: IB SMA v0.2
Nov 19 22:09:18 topspin-cc notifier.x[770]: [INFO]: Notifier v0.02
Nov 19 22:09:19 topspin-cc watchd_mgr.x[789]: [INFO]: Watchdog Manager v2.00
Nov 19 22:09:20 topspin-cc ip_mgr.x[812]: [INFO]: IP Manager v0.04
Nov 19 22:09:20 topspin-cc watchd_mgr.x[789]: [INFO]: process 4 is not up yet
Nov 19 22:09:21 topspin-cc fib_mgr.x[821]: [INFO]: FIB Manager v0.02
Nov 19 22:09:21 topspin-cc ip_mgr.x[812]: [INFO]: connected to watchd service, s
ent first pulse.
Nov 19 22:09:22 topspin-cc ib_mgr.x[843]: [INFO]: IB Manager v0.2
Nov 19 22:09:22 topspin-cc notifier.x[770]: [INFO]: connected to watchd service,
sent first pulse.
Nov 19 22:09:23 topspin-cc fib_mgr.x[821]: [INFO]: connected to watchd service,
sent first pulse.
Nov 19 22:09:23 topspin-cc srp_mgr.x[864]: [INFO]: SRP Manager v1.13, chassis-id
0x4000
Nov 19 22:09:23 topspin-cc ib_mgr.x[843]: [INFO]: connected to watchd service, s
ent first pulse.
Nov 19 22:09:23 topspin-cc srp_mgr.x[864]: [INFO]: Initializing DM.....
Nov 19 22:09:23 topspin-cc srp_mgr.x[864]: [INFO]: srpmDm initialized
<output truncated>

```

Show Most Recent Log Entries

To display the tail-end of the active log-file and display new log entries as they occur:

1. Set terminal parameters, if you have not already done so. Refer to [“Setting Terminal Parameters” on page 22](#).
It is recommended you set the terminal page length to 0 when using the end argument. Otherwise, you will have to keep pressing the <space> bar to continue each time the maximum display length is reached.
2. Enter the **show logging end** command.

```

Topspin-90# show logging end
Jan 1 00:02:30 igr-cc port_mgr.x[534]: [INFO]: port up - port=5/0,
type=ipGateway
Jan 1 00:02:34 s9 fc_portagent.x[448]: [port_config.c:149]: set AdminSpeed -
port=9/2, speed=5

```

The **show logging** command with the **end** argument locks the terminal window to display log. When using the **show logging** command without the **end** argument, log text is displayed a page at a time, as set by the terminal `length` parameter, much like the UNIX `tail -f` command.

3. Enter <Ctrl> C to stop displaying log entries and resume control of the terminal screen.

Show Details of a Specific Log

The **more** command displays the contents of a specified log file, including the active log-file.


1. Enter **more syslog:file name**

```
Topspin-360# more syslog:ts_log
Nov 19 22:08:49 topspin-cc kernel: THH kernel module initialized successfully
Nov 19 22:09:17 topspin-cc ts_sma.x[746]: [INFO]: IB SMA v0.2
Nov 19 22:09:18 topspin-cc notifier.x[770]: [INFO]: Notifier v0.02
Nov 19 22:09:19 topspin-cc watchd_mgr.x[789]: [INFO]: Watchdog Manager v2.00
Nov 19 22:09:20 topspin-cc ip_mgr.x[812]: [INFO]: IP Manager v0.04
Nov 19 22:09:20 topspin-cc watchd_mgr.x[789]: [INFO]: process 4 is not up yet
Nov 19 22:09:21 topspin-cc fib_mgr.x[821]: [INFO]: FIB Manager v0.02
Nov 19 22:09:21 topspin-cc ip_mgr.x[812]: [INFO]: connected to watchd service, sent first pulse.
Nov 19 22:09:22 topspin-cc ib_mgr.x[843]: [INFO]: IB Manager v0.2
Nov 19 22:09:22 topspin-cc notifier.x[770]: [INFO]: connected to watchd service, sent first pulse.
Nov 19 22:09:23 topspin-cc fib_mgr.x[821]: [INFO]: connected to watchd service, sent first pulse.
Nov 19 22:09:23 topspin-cc srp_mgr.x[864]: [INFO]: SRP Manager v1.13, chassis-id 0x4000
Nov 19 22:09:23 topspin-cc ib_mgr.x[843]: [INFO]: connected to watchd service, sent first pulse.
Nov 19 22:09:23 topspin-cc srp_mgr.x[864]: [INFO]: Initializing DM.....
Nov 19 22:09:23 topspin-cc srp_mgr.x[864]: [INFO]: srpmDm initialized
Nov 19 22:09:23 topspin-cc srp_mgr.x[864]: [INFO]: Initializing DM complete.
Nov 19 22:09:23 topspin-cc watchd_mgr.x[789]: [INFO]: process started: app=ip-mgr, pid=812, fd=11
Nov 19 22:09:23 topspin-cc watchd_mgr.x[789]: [INFO]: process 4 is not up yet
Nov 19 22:09:24 topspin-cc chassis_mgr.x[875]: [INFO]: Chassis Manager v2.0
Nov 19 22:09:24 topspin-cc srp_mgr.x[864]: [INFO]: connected to watchd service, sent first pulse.
```

Viewing a Log File Through the GUI

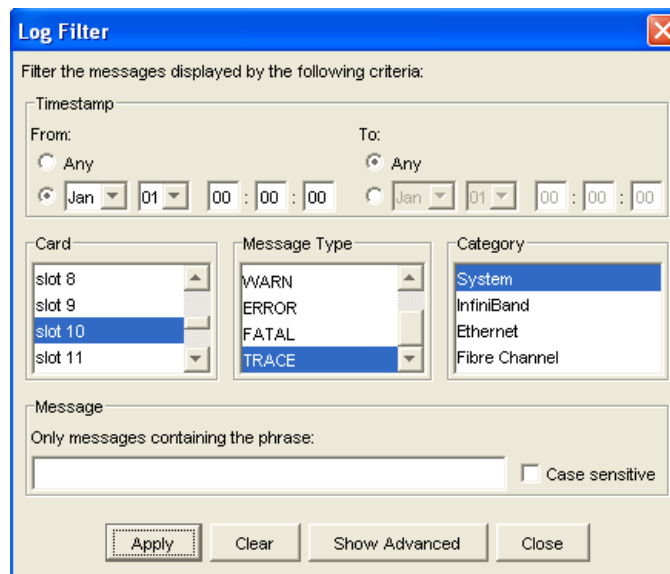
1. Launch Element Manager, if you have not already done so.
2. Select **Health --> Log Viewer**.

The Log Viewer window appears.



Filtering Logs

1. Launch Element Manager, if you have not already done so.
2. Select **Health --> Log Viewer**.
The Log Viewer window appears.
3. Click the **filter** button at the top of the window.
The Log Filter window appears.

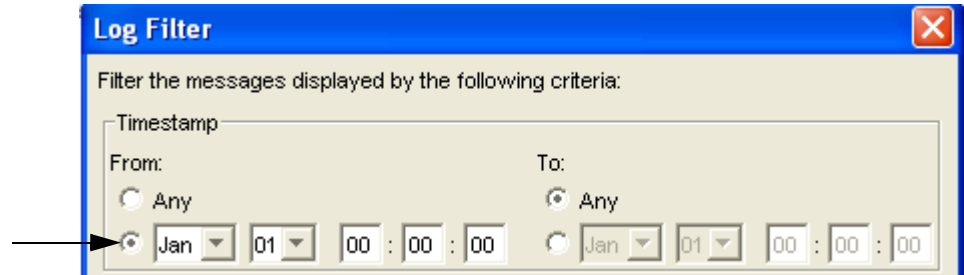


4. Determine the type of filter that you want to use:
 - [“Filter Logs by Time” on page 116](#)
 - [“Filter Logs by Card Slot” on page 116](#)
 - [“Filter Logs by Message Type” on page 116](#)

- “Filter Logs by Category” on page 117
- “Filter Logs by Text String” on page 117

Filter Logs by Time


1. Follow the steps in “Filtering Logs” on page 115.
2. Click the **From** date radio button.



3. Select the **Month** and **Day** from the drop-down menus.
4. Enter the Hour, Minute, and second in the fields that follow the date.
5. Click the **To** radio button.
6. Select the **Month** and **Day** from the drop-down menus.
7. Enter the Hour, Minute, and second in the fields that follow the date.
8. Click the **Apply** button.

Filter Logs by Card Slot


1. Follow the steps in “Filtering Logs” on page 115.
2. Select the Slot from which you want to obtain logs from the **Card** column.
Select multiple cards by holding down the <Ctrl> key while you left-click a slot.



3. Click the **Apply** button.

Filter Logs by Message Type


1. Follow the steps in “Filtering Logs” on page 115.
2. Click on the lowest level of message that you want to be captured in the log from the **Message Type** column.
Refer to “About Message Types” on page 111 for more information.



3. Click the **Apply** button.

Filter Logs by Category


1. Follow the steps in “[Filtering Logs](#)” on page 115.
2. Click on the technology for which you want to capture logs from the **Category** column.



3. Click the **Apply** button.

Filter Logs by Text String

1. Follow the steps in “[Filtering Logs](#)” on page 115.
2. Enter the text string that you want to use as your log filter into the **Message** field.



3. (Optional) Check the **Case Sensitive** box to further restrict the events that are captured by the log.
4. Click the **Apply** button.

Configuring Remote Logging

Logs can be configured to be saved to a remote host.

1. Enter the **logging** command and the ip address of the remote host where the log files will be saved.

```
Topspin-360> enable
Topspin-360# configure
Topspin-360(config)# logging 10.3.102.60
```


Viewing the IB Network Through the GUI

This chapter provides the following information:

- [“About the Device Manager \(DM\)” on page 119](#)
- [“Display the Device Manager” on page 119](#)
- [“About the Topology View” on page 121](#)
- [“Display the InfiniBand Topology” on page 122](#)
- [“View the Internal Chassis Topology” on page 127](#)
- [“View Subnet Manager Details” on page 129](#)

About the Device Manager (DM)

Use the Device Manager to view the InfiniBand I/O units, I/O Controller information, and I/O Controller services.

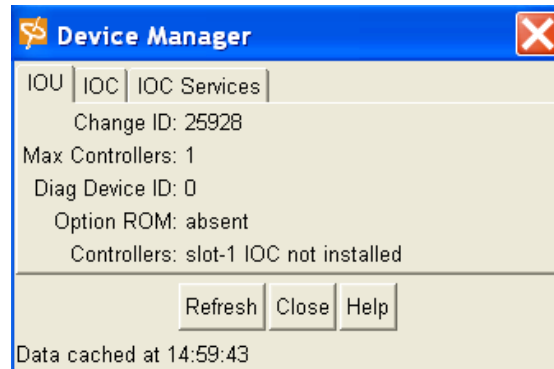
The Device Manager is available through the Element Manager GUI, the Chassis Manager GUI, as well as the CLI.

Display the Device Manager

View I/O Unit Information

1. Launch Element Manager, if you have not already done so.
2. Select **InfiniBand --> DM**

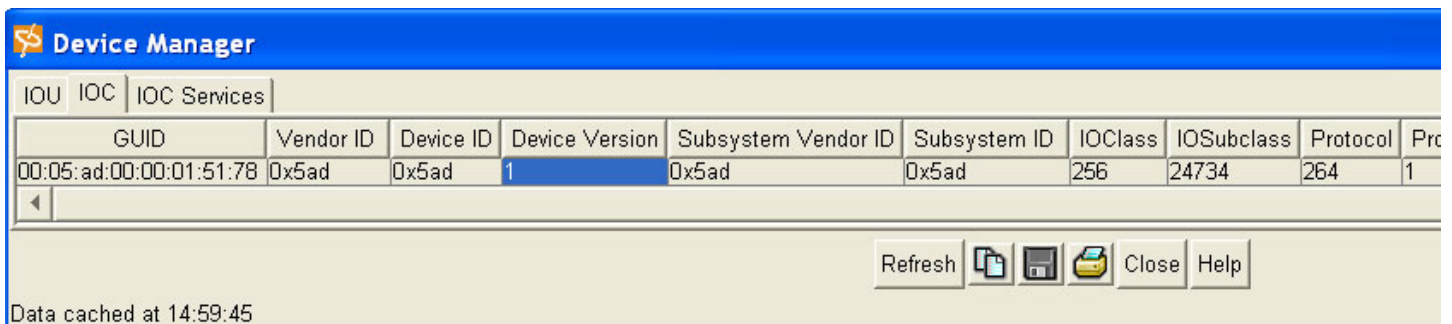
The Device Manager window opens.



3. View the cumulative number of changes to the Controller list since the device last booted in the **Change ID** field.
4. View the maximum number of controllers that your device can support in the **Max Controllers** field.
5. View the **Diag Device ID** field. The number 1 is displayed if diagnostics can provide IOC details; otherwise the field displays 0.
6. View the **Option ROM** field to determine the presence or absence of Option ROM.
7. View the **Controllers** field for a list of each slot on your Server Switch that can potentially contain a controller, and to identify whether or not a controller resides in that slot.

View I/O Controller Units

1. Launch Element Manager, if you have not already done so.
2. Select **InfiniBand --> DM**
The Device Manager window opens.
3. Click the **IOC** Tab.



4. View the Global Unique Identifier for the I/O Controller in the **GUID** field.
5. View the organization-Unique Identifier (OUI) of the vendor in the **Vendor ID** field.
6. View the vendor-assigned IOC device identifier in the **Device ID** field.
7. View the vendor-assigned subsystem identifier of the vendor in the **Subsystem Vendor ID** field.
8. View the vendor-assigned subsystem identifier in the **Subsystem ID** field.
9. View the I/O class that is supported by the I/O controller in the **IOClass** field.
10. View the subclass of the I/O class protocol of the IOC in the **IOSubclass** field.
11. View the standard protocol definition that is supported by the I/O Controller in the **Protocol** field.
12. View the protocol version that is supported by the I/O Controller in the **Protocol Version** field.
13. View the maximum number of messages that the send message queue supports in the **Send Message Queue Depth** field.

14. View the maximum depth of the per-channel RDMA Read Queue in the **RDMA Read Queue Depth** field.
15. View the maximum size, in bytes, of send messages in the **Send Message Size** field.
16. View the maximum size, in bytes, of the outbound RDMA transfers that the IOC initiates in the **RDMA Transfer Size** field.
17. View the integer value (from 8 cumulative bits) between 1 and 255 that represents the operation type(s) that the IOC supports in the **Controller Op Cap** field.

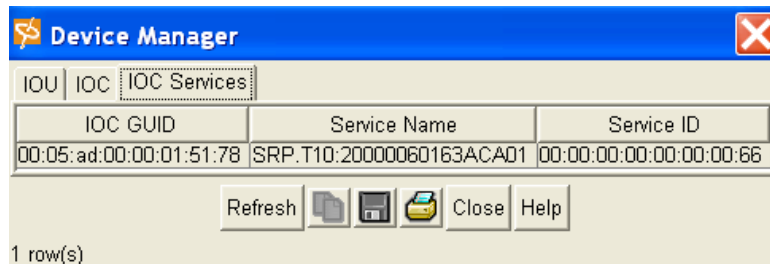
The entries can be read as the following:

- bit 0: ST; Send Messages To IOCs
- bit 1: SF; Send Messages From IOCs
- bit 2: RT; RDMA Read Requests To IOCs
- bit 3: RF; RDMA Read Requests From IOCs
- bit 4: WT; RDMA Write Requests To IOCs
- bit 5: WF; RDMA Write Requests From IOCs
- bit 6: AT; Atomic Operations To IOCs
- bit 7: AF; Atomic Operations From IOCs

18. View the number of services that the IOC provides in the **Services Entries** field.

View I/O Controller Units Services

1. Launch Element Manager, if you have not already done so.
2. Select **InfiniBand --> DM**
The Device Manager window opens.
3. Click the IOC Services Tab.



4. View the Global Unique Identifier for the I/O node that provides the service in the **IOC Services** field.
5. View the ASCII identifier of the service in the **Service Name** field.
6. View the numeric identifier that nodes use to call the service in the **Service ID** field.

About the Topology View

The Topology view is available through the Element Manager GUI, and is an easy way to view all of the elements in an InfiniBand network.

The Topology view is available by selecting **InfiniBand --> Topology**.

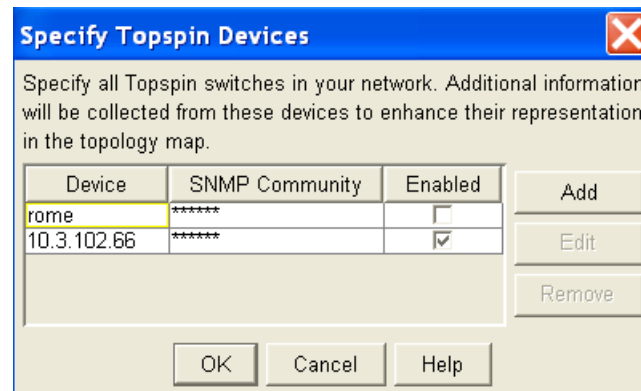
Display the InfiniBand Topology

View the Topology

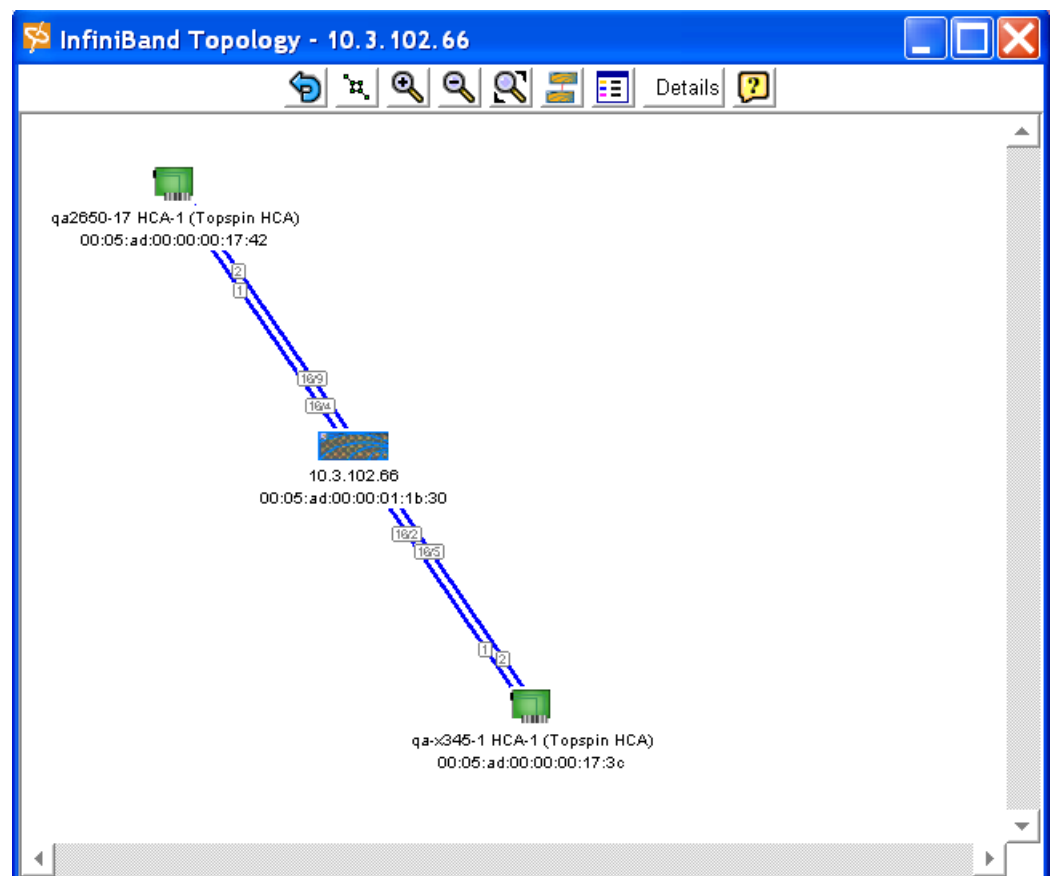
1. Launch Element Manager, if you have not already done so.
2. Select **InfiniBand --> Topology**.

The Specify Devices dialog box opens, if this is your first time viewing the Topology.

To add a device, refer to [“Add an Attached Device to the Topology View”](#) on page 126.



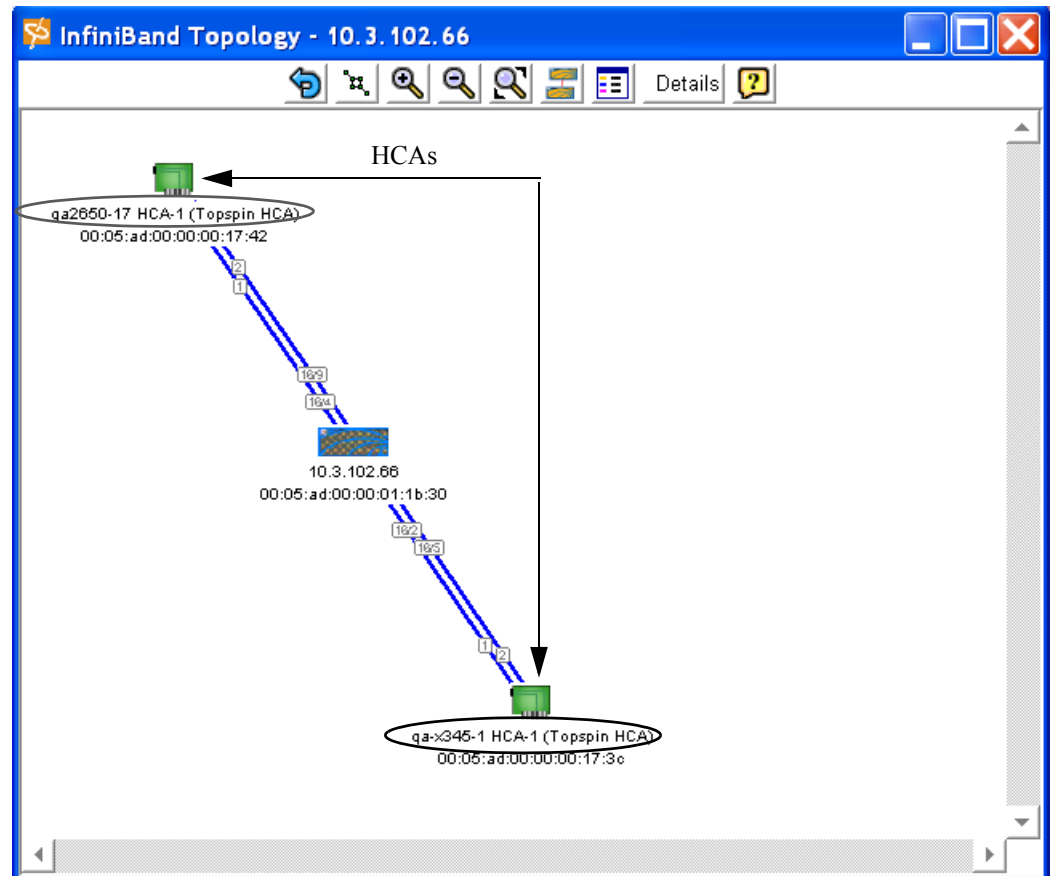
3. Click the **OK** button.
The Topology view appears.



View the Name of an HCA

To easily view the name of a Host Channel Adapter, use the Topology view.


1. Open the Topology view, if you have not already done so. Refer to [“View the Topology”](#) on page 122.
2. View the name of the HCA that is displayed with the HCA icons.



View the GUID of an HCA

To easily associate a Host Channel Adapter with the Global Unique Identifier (GUID), use the Topology view.

1. Open the Topology view, if you have not already done so. Refer to [“View the Topology” on page 122](#).
2. View the GUID information that is displayed with the HCA icons.



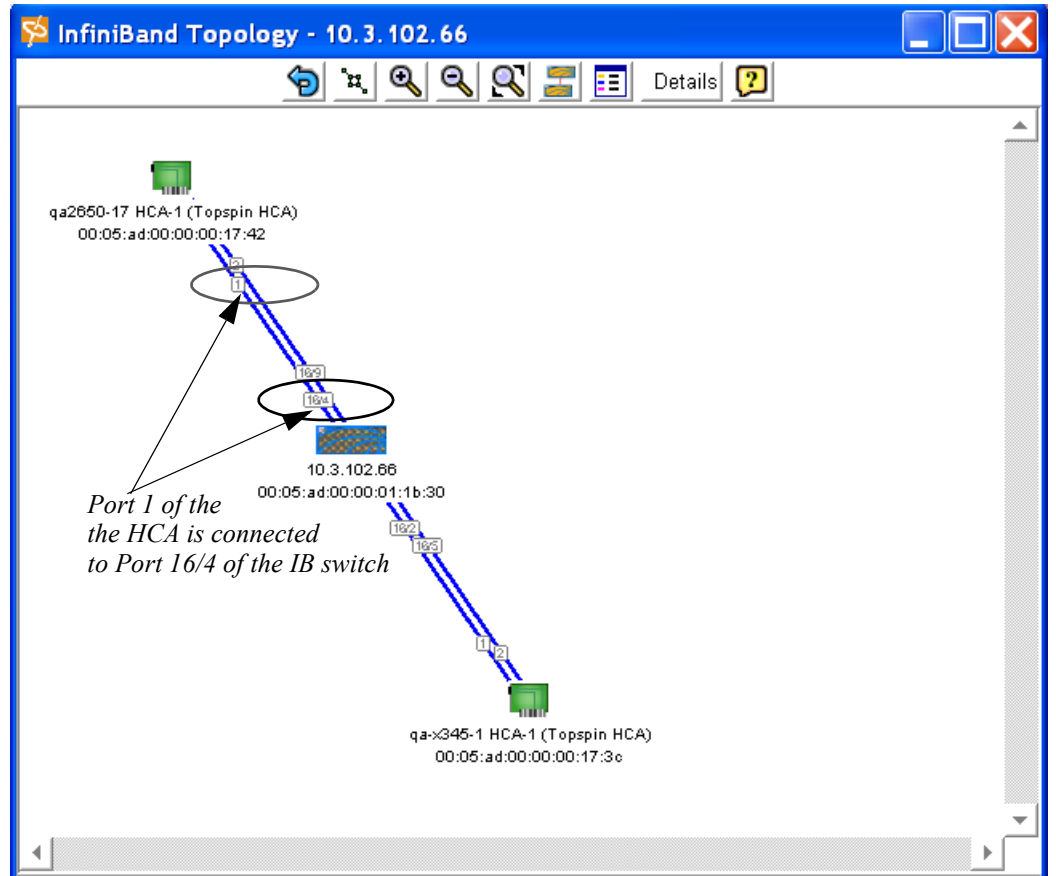
Determine Which HCA Port is Connected to an IB Port

To easily determine which Host Channel Adapter (HCA) port is attached to a specific InfiniBand (IB) port, use the Topology view.

Viewing the ports in this way prevents you from having to view the physical hardware to determine which ports are connected.

1. Open the Topology view, if you have not already done so. Refer to [“View the Topology” on page 122](#).
2. View the HCA port and follow the connection to the IB switch.
 - The port number of the HCA is labelled on the link close to the HCA icon.

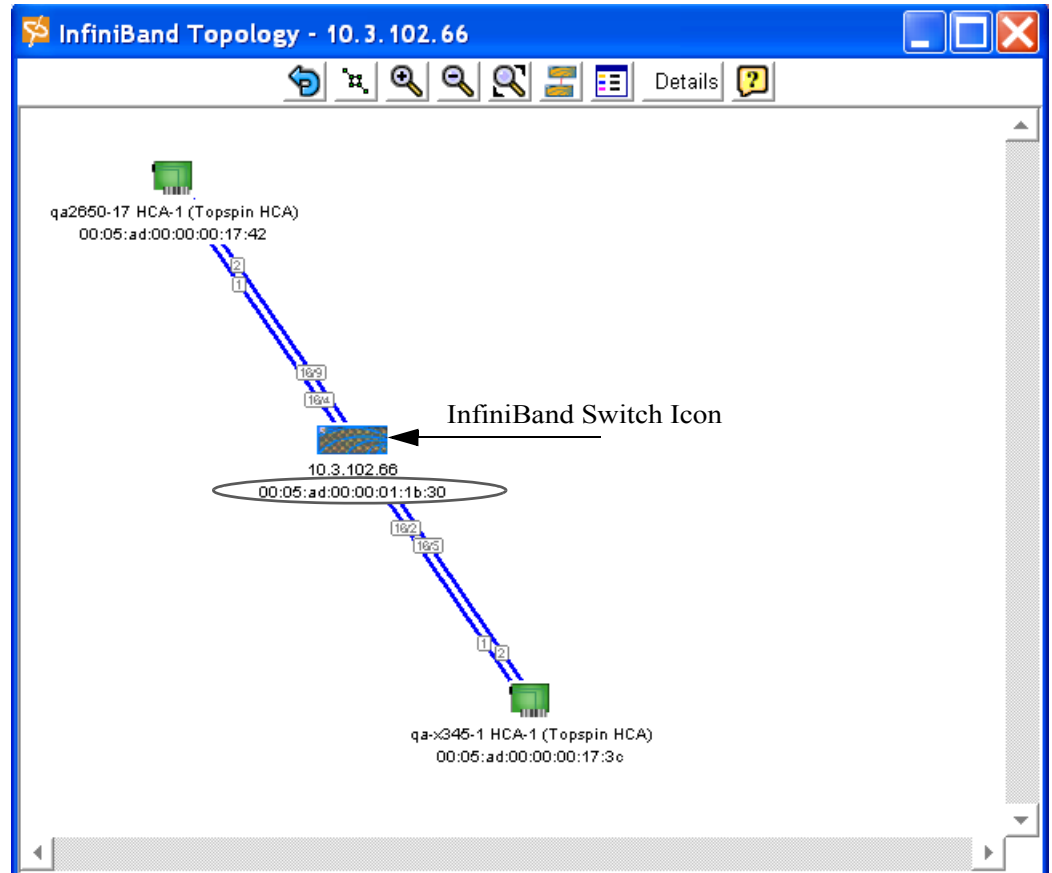
- The slot/port number of the IB switch is labelled close to the IB switch icon.



View the GUID of an IB Switch

To easily associate an InfiniBand (IB) switch with the Global Unique Identifier (GUID), use the Topology view.


1. Open the Topology view, if you have not already done so. Refer to “[View the Topology](#)” on page 122.
2. View the switch GUID information that is displayed with the InfiniBand switch icon(s).



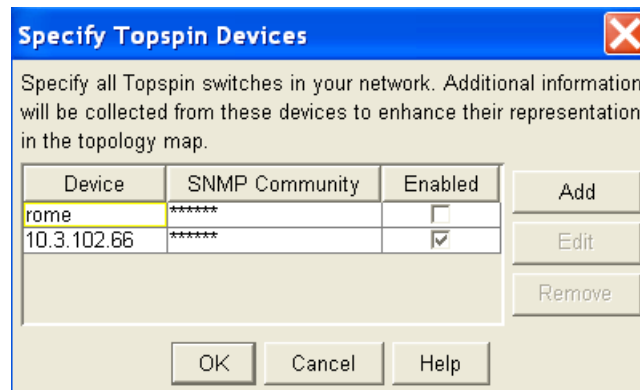
Add an Attached Device to the Topology View

When you first open the Topology view in a given session, the Element Manager opens a dialog box that allows you to include any new devices in your Topology.


However, if you want to add a new device to the Topology after you have already opened the Topology view, perform the following steps:

1. Select **InfiniBand --> Topology**.
The Topology view appears (unless this is a new session). If this is a new session, the Specify Devices dialog box appears. Skip step 2.
2. Click the **Specify Devices** icon in the Topology view. 

The Specify Devices dialog box appears.



3. Click the **Add** button to add another connected InfiniBand device to the Topology view. The Add dialog box appears.

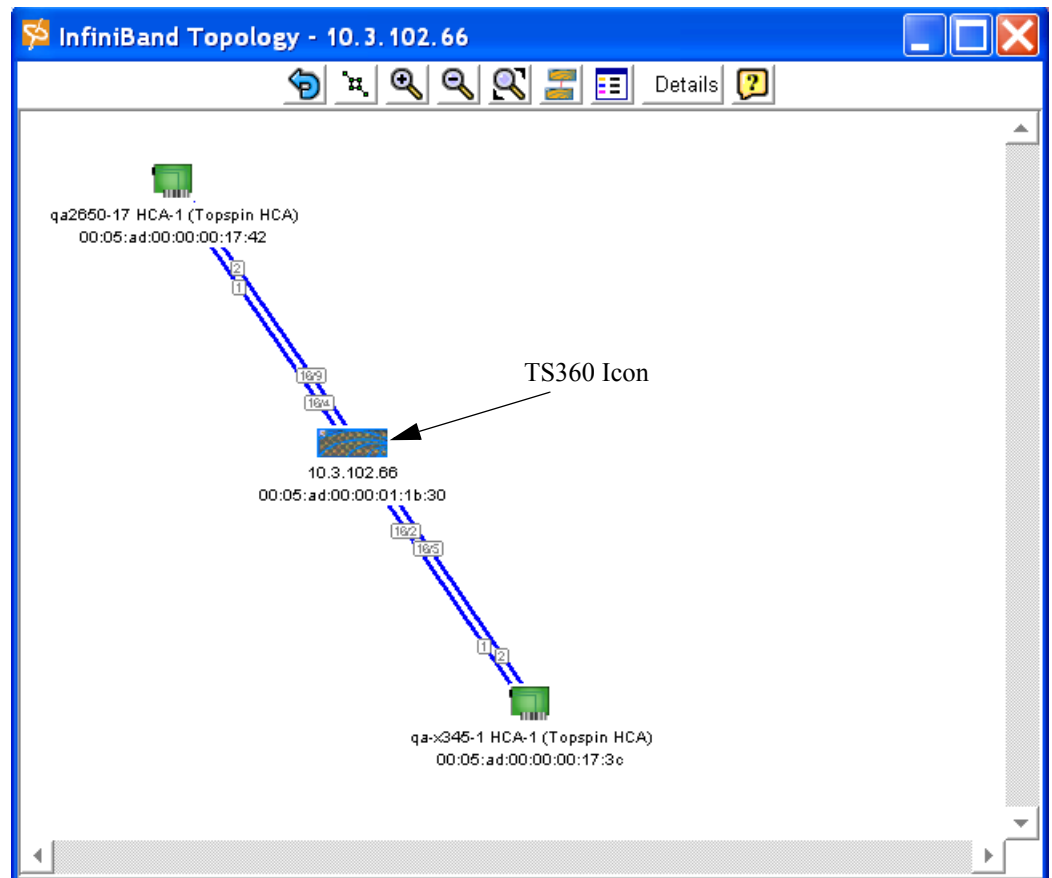


4. Enter the DNS name or the Management-Ethernet port IP address of the InfiniBand device in the **Device Name or IP Address** field.
5. Enter the SNMP Community string for the InfiniBand device.
6. Check or uncheck the **Device is Enabled** box.
If the connected chassis is running and you want it included in the Topology Manager, check the Device is Enabled toggle. You can also enable this later by going back to the Specify Device window.
7. Click the **Add** button.
8. The connected device is added to the Topology view.

View the Internal Chassis Topology


1. Select **InfiniBand --> Topology**.
The Specify Devices dialog box appears (if this is a new Element Manager session).
2. Click the **OK** button.

The external Topology view appears.



3. Double-click the InfiniBand switch icon. The icon will appear different, depending on the type of IB switch you are using in your network.

The **Internal InfiniBand Topology** appears.



4. View the connections between InfiniBand nodes and Ethernet or Fibre Channel Gateways.
5. View the descriptions and node GUIDs of the InfiniBand nodes and any gateways.

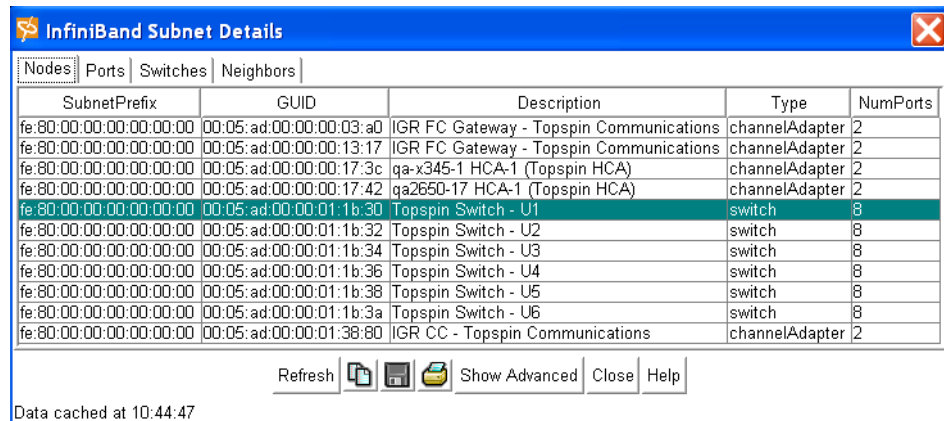
View Subnet Manager Details

In addition to initializing and maintaining the InfiniBand fabric, the subnet manager also communicates with subnet agents to track changes to the InfiniBand topology as they occur. The information recorded by the subnet manager can be viewed in a table format through the Element Manager.

View Basic Node Information

1. Launch the Element Manager, if you have not already done so.
1. Select **InfiniBand** --> **Topology**.
The Specify Devices dialog box appears (if this is a new Element Manager session).
2. Click the **OK** button, if it appears.
The Topology view appears.
3. Click the **Details** button at the top of the Topology view.

- The **InfiniBand Subnet Details** window appears.



SubnetPrefix	GUID	Description	Type	NumPorts
fe:80:00:00:00:00:00:00	00:05:ad:00:00:00:03:a0	IGR FC Gateway - Topspin Communications	channelAdapter	2
fe:80:00:00:00:00:00:00	00:05:ad:00:00:00:13:17	IGR FC Gateway - Topspin Communications	channelAdapter	2
fe:80:00:00:00:00:00:00	00:05:ad:00:00:00:17:3c	qa-x345-1 HCA-1 (Topspin HCA)	channelAdapter	2
fe:80:00:00:00:00:00:00	00:05:ad:00:00:00:17:42	qa2650-17 HCA-1 (Topspin HCA)	channelAdapter	2
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:30	Topspin Switch - U1	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:32	Topspin Switch - U2	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:34	Topspin Switch - U3	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:36	Topspin Switch - U4	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:38	Topspin Switch - U5	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:3a	Topspin Switch - U6	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:38:80	IGR CC - Topspin Communications	channelAdapter	2

Refresh [Print] [Save] [Export] Show Advanced Close Help

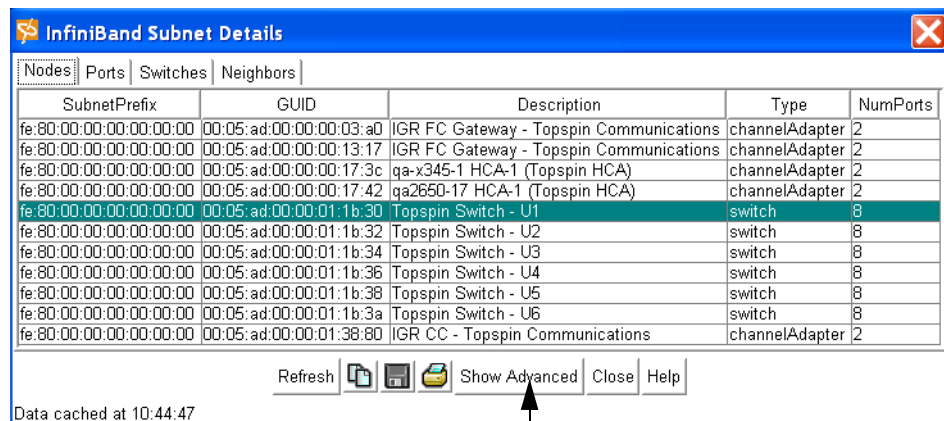
Data cached at 10:44:47

The Node tab is selected by default.

- View the Subnet Prefix of the node. The **SubnetPrefix** field identifies the InfiniBand subnet to which the node belongs.
- View the Global Unique Identifier (**GUID**) of the node.
- View a text string that describes the node is the **Description** field.
- View the kind of node that is being managed in the **Type** field. The value is *channelAdapter*, *switch*, *route*, or *error*. The *error* value indicates an unknown type.
- View the number of physical ports available on the node in the **NumPorts** field.
- Continue to “[View Advanced Node Information](#)” on page 130 for more information.

View Advanced Node Information

- Follow the steps in “[View Basic Node Information](#)” on page 129.
- Click the **Show Advanced** button in the InfiniBand Subnet Details window.

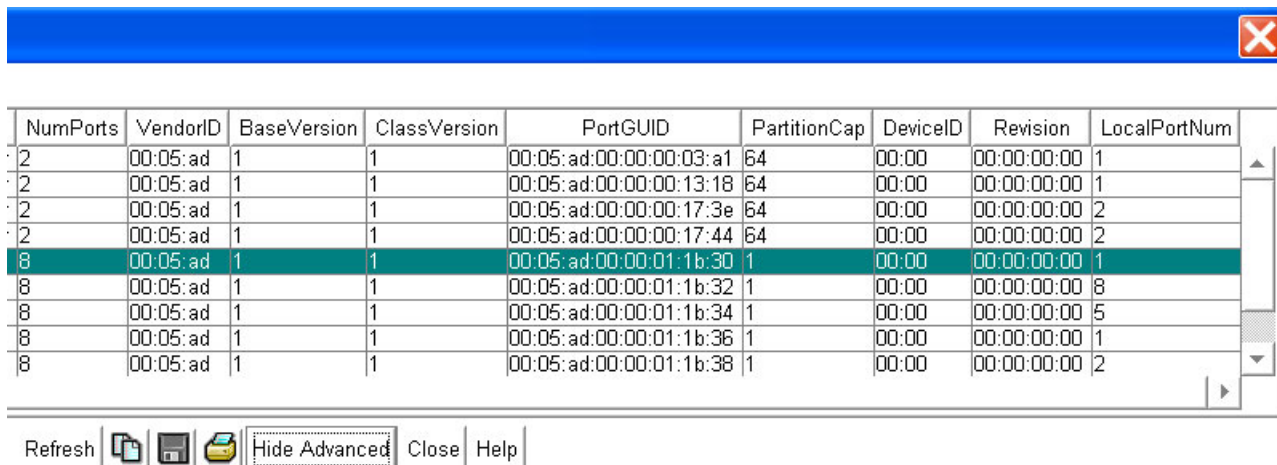


SubnetPrefix	GUID	Description	Type	NumPorts
fe:80:00:00:00:00:00:00	00:05:ad:00:00:00:03:a0	IGR FC Gateway - Topspin Communications	channelAdapter	2
fe:80:00:00:00:00:00:00	00:05:ad:00:00:00:13:17	IGR FC Gateway - Topspin Communications	channelAdapter	2
fe:80:00:00:00:00:00:00	00:05:ad:00:00:00:17:3c	qa-x345-1 HCA-1 (Topspin HCA)	channelAdapter	2
fe:80:00:00:00:00:00:00	00:05:ad:00:00:00:17:42	qa2650-17 HCA-1 (Topspin HCA)	channelAdapter	2
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:30	Topspin Switch - U1	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:32	Topspin Switch - U2	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:34	Topspin Switch - U3	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:36	Topspin Switch - U4	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:38	Topspin Switch - U5	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:1b:3a	Topspin Switch - U6	switch	8
fe:80:00:00:00:00:00:00	00:05:ad:00:00:01:38:80	IGR CC - Topspin Communications	channelAdapter	2

Refresh [Print] [Save] [Export] Show Advanced Close Help

Data cached at 10:44:47

An additional level of node information is added to the window.



NumPorts	VendorID	BaseVersion	ClassVersion	PortGUID	PartitionCap	DeviceID	Revision	LocalPortNum
2	00:05:ad	1	1	00:05:ad:00:00:00:03:a1	64	00:00	00:00:00:00	1
2	00:05:ad	1	1	00:05:ad:00:00:00:13:18	64	00:00	00:00:00:00	1
2	00:05:ad	1	1	00:05:ad:00:00:00:17:3e	64	00:00	00:00:00:00	2
2	00:05:ad	1	1	00:05:ad:00:00:00:17:44	64	00:00	00:00:00:00	2
8	00:05:ad	1	1	00:05:ad:00:00:01:1b:30	1	00:00	00:00:00:00	1
8	00:05:ad	1	1	00:05:ad:00:00:01:1b:32	1	00:00	00:00:00:00	8
8	00:05:ad	1	1	00:05:ad:00:00:01:1b:34	1	00:00	00:00:00:00	5
8	00:05:ad	1	1	00:05:ad:00:00:01:1b:36	1	00:00	00:00:00:00	1
8	00:05:ad	1	1	00:05:ad:00:00:01:1b:38	1	00:00	00:00:00:00	2

3. View the Device vendor ID in the **DeviceID** field. The value will be the same for all ports on the node.
4. View the supported base management datagram (MAD) version in the **BaseVersion** field. This field indicates that the channel adapter, switch, or router supports up to and including this version.
5. View the supported base management datagram (MAD) class format in the **ClassVersion** field. This field indicates that the channel adapter, switch, or router supports up to and including this version.
6. View the GUID of a port on the node in the **PortGUID** field. A port within a node can return the node GUID as its PortGUID if the port is an integral part of the node and is not field-replaceable (i.e., not swappable).
7. View the number of entries in the partition table for channel adapter, router, and the switch management port in the **PartitionCap** field. The value is the same for all ports on the node. This is set to at least 1 for all nodes including switches.
8. View the manufacturer-assigned device identification for the node in the **DeviceID** field.
9. View the manufacturer-assigned device revision. for the node in the **Revision** field.
10. View the link port number on which a subnet management packet (SMP) came in the **LocalPortNum** field. The value is the same for all ports on the node.
11. Click the **Hide Advanced** or **Close** button.

View Basic Port Information

1. Select **InfiniBand --> Topology**.
The Specify Devices dialog box appears (if this is a new Element Manager session).
2. Click the **OK** button, if it appears.
The Topology view appears.
3. Click the **Details** button at the top of the Topology view.

- The **InfiniBand Subnet Details** window appears.

SubnetPrefix	GUID	Description	Type	NumPorts
fe:80:00:00:00:00:00	00:05:ad:00:00:03:a0	IGR FC Gateway - Topspin Communications	channelAdapter	2
fe:80:00:00:00:00:00	00:05:ad:00:00:13:17	IGR FC Gateway - Topspin Communications	channelAdapter	2
fe:80:00:00:00:00:00	00:05:ad:00:00:17:3c	ga-x345-1 HCA-1 (Topspin HCA)	channelAdapter	2
fe:80:00:00:00:00:00	00:05:ad:00:00:17:42	ga2650-17 HCA-1 (Topspin HCA)	channelAdapter	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	Topspin Switch - U1	switch	8
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	Topspin Switch - U2	switch	8
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	Topspin Switch - U3	switch	8
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:36	Topspin Switch - U4	switch	8
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:38	Topspin Switch - U5	switch	8
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:3a	Topspin Switch - U6	switch	8
fe:80:00:00:00:00:00	00:05:ad:00:00:01:38:80	IGR CC - Topspin Communications	channelAdapter	2

Refresh [Copy] [Save] [Print] Show Advanced Close Help

Data cached at 10:44:47

The Node tab is selected by default.

- Click the **Ports** tab.

SubnetPrefix	NodeGUID	Port	LID	State	LinkWidthActive
fe:80:00:00:00:00:00	00:05:ad:00:00:03:a0	1	2	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:13:17	1	3	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:17:3c	1	11	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:17:3c	2	12	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:17:42	1	193	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:17:42	2	194	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	0	4	active	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	1	0	active	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	2	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	3	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	4	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	5	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	6	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	7	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:30	8	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	0	5	active	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	1	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	2	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	3	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	4	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	5	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	6	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	7	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:32	8	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	0	6	active	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	1	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	2	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	3	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	4	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	5	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	6	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	7	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:34	8	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:36	0	7	active	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:36	1	0	active	2
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:36	2	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:36	3	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:36	4	0	down	1
fe:80:00:00:00:00:00	00:05:ad:00:00:01:1b:36	5	0	active	2

Refresh [Copy] [Save] [Print] Show Advanced Close Help

Data cached at 10:44:58

- View the Subnet Prefix of the node to which the port belongs. The **SubnetPrefix** field identifies the InfiniBand subnet to which the node belongs.
- View the Global Unique Identifier (**GUID**) of the node to which the port belongs.
- View the local port number on this node in the **Port** field.
- View the Local Identifier (**LID**) for the port in the **LID** field. The LID is assigned to a port by the subnet manager, and it used for directing packets *within* the subnet.

10. Determine whether or not the nodes can actually communicate, and view the state transition that has occurred in the **State** field.
A Transition is a port change from down to initialize, from initialize to down, from armed to down, or from active to down as a result of link state machine logic.
11. View the Active link width in the **LinkWidthActive** field. This field is used in conjunction with **LinkSpeedActive** (view Advanced section) to determine the link rate between two nodes.
The values are: 1 (1x), 2 (4x), or 8 (12x).
12. Continue to “[View Basic Port Information](#)” on page 131 for more information.

View Advanced Port Information

1. Follow the steps in “[View Basic Port Information](#)” on page 131.
2. Click the **Show Advanced** button.

An additional level of information is added to the window.

InfiniBand Subnet Details											
Nodes Ports Switches Neighbors											
MKey	GIDPrefix	MasterSmLID	CapMask	DiagCode	MKeyLeasePeriod	LinkWidthEnabled	LinkWidthSupported	LinkSpeedSupported	PhyState	LinkDownDefStat	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	1	00:10:0a:68	10:29	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	1	00:10:0a:68	10:29	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	1	00:50:0a:68	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	1	00:50:0a:68	10:29	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	1	00:50:0a:68	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	1	00:50:0a:68	10:29	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	00:00:00:00:00:00:00:00	1	00:00:0a:48	10:29	10985	3	3	1	linkup	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	00:00:00:00:00:00:00:00	1	00:00:0a:48	10:29	10985	3	3	1	linkup	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	00:00:00:00:00:00:00:00	1	00:00:0a:48	10:29	10985	3	3	1	linkup	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	00:00:00:00:00:00:00:00	1	00:00:0a:48	10:29	10985	3	3	1	linkup	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	
00:00:00:00:00:00:00:00	fe:80:00:00:00:00:00:00	0	00:00:00:00	00:00	15	3	3	1	noStateChange	polling	

3. View the management key for the port in the **MKey** field. The management key is used to authenticate a sender to a receiver.
4. View the Global Identifier (GID) prefix in the **GIDPrefix** field. The GID prefix is assigned by the subnet manager, based upon the port router and the rules for local identifiers.
5. View the base Local Identifier (LID) of the subnet manager that is managing this port in the **MasterSmLID** field. The LID is assigned to a port by the subnet manager, and it used for directing packets *within* the subnet.
6. View a bitmask that specifies the supported capabilities of the port in the **CapMask** field. A bit value of 1 (one) indicates a supported capability. The bits are:
 - 0 , 11 -15 , 18 , 21 -31 (Reserved is always 0).
7. View a 16-bit diagnostic code in the **DiagCode** field. For all ports, all bits set to zero means the port status is good. Any non-zero value means there may be error conditions.

8. View the initial value of the lease-period timer in seconds in the **MKeyLeasePeriod** field. Refer to section 14.2.4, Management Key, *InfiniBand® Architecture*, Vol. 1, Release 1.0, for more information.
9. View the enabled link width in the **LinkWidthEnabled** field. The value is an integer that indicates the enabled link-width sets for this port. The value may be:
 - 0 (no state change),
 - 1 (1x)
 - 2 (4x)
 - 3 (1x or 4x)
 - 8 (12x)
 - 9 (1x or 12x)
 - 10 (4x or 12x)
 - 11 (1x, 4x or 12x)
10. View the support link width for the port in the **LinkWidthSupported** field. The values are:
 - 1 = (1x)
 - 3 = (1x or 4x)
 - 11 (1x, 4x, or 12x)
11. View the support link speed for the port in the **LinkSpeedSupported** field. The value is:
 - 1 = 2.5 Gbps
12. View the actual state of the port in the **PhyState** field. This field determines that electricity is flowing between nodes and that they can hand-shake. The possible fields are:
 - noStateChange
 - sleeping
 - polling (this is the default state upon power-up)
 - disabled
 - portConfigurationTraining
 - linkup
 - linkErrorRecovery
13. View the state to which a down link will return by default in the **DefaultLinkDown** field. The possible fields are:
 - noStateChange
 - sleeping
 - polling
14. View the Management key protection bits for the port in the **MkeyProtBits** field. Refer to section 14.2.4.1, *Levels of Protection*, *InfiniBand® Architecture*, Vol. 1, Release 1.0, for more information.
15. View the Local identifier mask control (LMC) for multipath support in the **LMC** field. An LMC is assigned to each channel adapter and router port on the subnet. It provides multiple virtual ports within a single physical port. The value of the LMC specifies the number of path bits in the LID. A value of 0 (zero) indicates one LID is allowed on this port.
16. View the speed of an active link in the **LinkSpeedActive** field. The value is 1 (2.5Gbps).
17. View the maximum speed the link is capable of handling in the **LinkSpeedEnabled** field. The possible fields are:
 - 0 = (No state change)
 - 1 = (2.5Gbps)
 - 3 = (value derived from **LinkSpeedSupported**)

18. View the active maximum transmission unit enabled on this port for transmit in the **NeighborMTU** field. Check the **MTUCap** value at both ends of every link and use the lesser speed.
19. View the administrative service level required for this port to send a non-subnet management packet (SMP) message to the subnet manager in the **MasterSmSL** field.
20. View the maximum range of data virtual lanes supported by this port in the **VLCap** field. The value are:
 - v10
 - v10ToV11
 - v10ToV13
 - v10ToV17
 - v10ToV114
21. See the **VLArbHighCap** field to view the maximum high-priority limit on the number of bytes allowed for transmitting high-priority packets when both ends of a link operate with multiple data virtual-lanes. This field is used with the virtual-lane arbitration table and specified as a VL/Weight pair.
22. See the **VLArbLowCap** field to view the lowest arbitration value allowed by the arbiter in determining the next packet in a set of packets to send across the link.
23. View the **MTUCap** field. This field is used in conjunction with **NeighborMTU** to determine the maximum transmission size supported on this port. The lesser of MTUCap and NeighborMTU determines the actual MTU used.
The values are: mtu256, mtu512, mtu1024, mtu2048, or mtu4096.
24. View the number of sequentially dropped packets at which the port enters a VLStalled state in the **VLStallCount** field. The virtual lane exits the VLStalled state (8 * HLL) units after entering it. See section 18.2.5.4, Transmitter Queuing , *InfiniBand® Architecture, Vol. 1, Release 1.0*, for a description of HLL.
25. View the maximum duration allowed to packets at the head of a virtual-lane queue in the **HOQLife** field. This field is used with **VLStallCount** to determine the outgoing packets to discard.
26. View the administrative limit for the number of virtual lanes allowed to the link in the **OperVL** field. The values are: v10, v10ToV11, v10ToV13, v10ToV17, or v10ToV114.
27. View the Boolean value that indicates whether or not to support optional partition enforcement for the packets received by this port in the **InPartEnforce** field.
28. View the Boolean value that indicates whether or not to support optional partition enforcement for the packets transmitted by this port in the **OutPartEnforce** field.
29. View the Boolean value that indicates whether or not this port supports optional raw packet enforcement for the raw packets received by this port in the **InFilterRawPktEnf** field.
30. View the Boolean value that indicates whether or not this port supports optional raw packet enforcement for the raw packets transmitted by this port in the **OutFilterRawPktEnf** field.
31. View the number of subnet management packets (SMPs) that have been received on this port with *invalid* M_Keys since initial power up or the last reset in the **MKeyViolation** field.
32. View the number of subnet management packets (SMPs) that have been received on this port with *invalid* P_Keys since initial power up or the last reset in the **PKeyViolation** field.
33. View the number of subnet management packets (SMPs) that have been received on this port with *invalid* Q_Keys since initial power up or the last reset in the **QKeyViolation** field.
34. View the number of Global Unique Identifiers (GUID) entries allowed for this port in the port table in the **GUIDCap** field. Any entries that exceed this value are ignored on write and read back as zero.
35. View the maximum propagation delay allowed for this port to reach any other port in the subnet in the **SubnetTimeout** field. This value also affects the maximum rate at which traps can be sent from this port. Delay is affected by switch configuration. This parameter, along with **RespTime**, may be

used by requestors to determine the interval to wait for a response to a request before taking other action. Duration is calculated as $(4.096 \text{ ms} * 2^{\text{SubnetTimeout}})$.

36. View the maximum time allowed between the port reception of a subnet management packet and the transmission of the associated response in the **RespTime** field.
37. View the **LocalOverrunError** field to view the threshold at which the count of buffer overruns, across consecutive flow-control update periods, and will result in an overrun error.
38. View the **LocalPhyError** field to view the threshold at which ICRC, VCRC, FCCRC, and all physical errors result in an entry into the BAD PACKET or BAD PACKET DISCARD states of the local packet receiver.
39. Click the **Hide Advanced** or **Close** button.

Monitoring and Reporting Through the GUI

This chapter gives an overview of the following:

- [“About Analyzing Network Data” on page 137](#)
- [“About Tabular Formats” on page 138](#)
- [“About Graph Formats” on page 138](#)
- [“Creating a Data Analysis Table” on page 140](#)
- [“Creating a Data Analysis Graph” on page 143](#)
- [“About SNMP Traps” on page 146](#)
- [“Configuring SNMP Settings” on page 147](#)

About Analyzing Network Data

The GUI is a convenient tool for tracking and analyzing network activity across interface ports and cards. Statistical data, such as the number of late collisions and received datagrams, are automatically tallied and updated.

Benefits

The Element Manager simplifies the command entry process. You may perform complex chassis configuration procedures in the Element Manager with a few quick mouse-clicks.

Using standard point-and-click methodology, the Element Manager is used to

- track network traffic and changes
- report the health and link activity of the InfiniBand system

Data Captured

The data displayed depends upon the selected cards and ports. General interface activity, such as the number of octets, packets, and transmission errors, is provided for all port types.

About Tabular Formats

Network data may be viewed in a tabular or graph format.

The tabular format is ideal for numbers crunching when you want to know precise activity statistics.

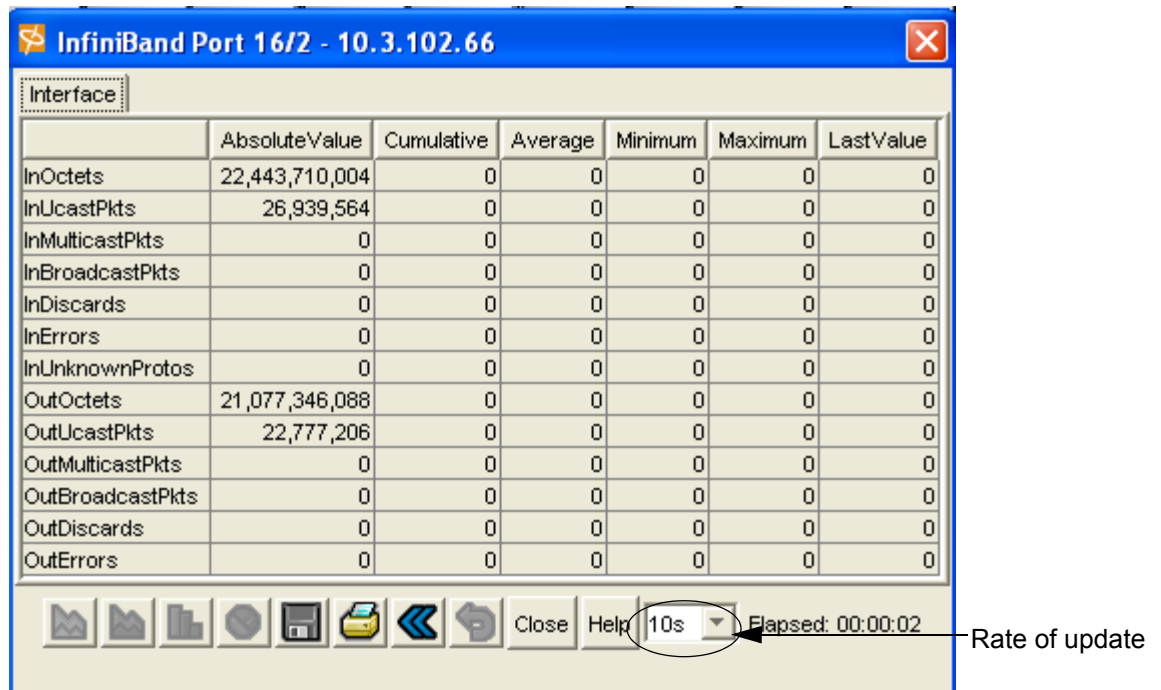


Figure 13-1: Single-port InfiniBand selected

Each row in the table is a parameter that is being logged, such as the number of multicast packets or discarded packets.

The objects in the first column identify what is being logged. These are the names of the objects being logged.

When only one port or card is selected, the remaining columns are counters that contain numeric data that is derived from either actual values or computed from actual values.

About Graph Formats

Network data may be viewed in a tabular or graph format.

The graph format provides a comparative view of the same data so you can evaluate differences at a glance.

Types of Graphs

The following types of graphs are available to visually depict network data:

- Pie
- Line

- Bar
- Area








Figure 13-2: In-Octets Statistics for a Single InfiniBand Port

The appropriate graph icons become enabled when a set of graph data is selected. You can click a graph type and modifier to indicate how to display that data.

Table 13-1: Graph Types and Descriptions

Graph Type	Associated Icon	Purpose
Pie		This graph represents counter data as a pie graph. Each wedge in the pie is a percentage of all the selected counters. Only a single data-set may be selected to display the data as a pie chart.
Line		This graph represents counter data as a lined graph over an arbitrary unit. When a single data-set is selected, such as the fields in one row or in one column, counter data is displayed as a line graph over time. It shows data as in the same format as the area graph but without fill colors.
Bar		This graph represents counter data as filled columns. When a single data-set is selected, such as the fields in one row or in one column, counter data is displayed as a filled columns over time.
Area		This graph represents counter data as a filled line graph over an arbitrary unit. When a single data-set is selected, such as the fields in one row or in one column, counter data is displayed as a filled line graph over time.




Visual Modifiers

Modifiers may be used singularly, or in conjunction with other modifiers, to change the orientation and scale of the graph displayed.

With the use of display modifiers, the data displayed can be:

- Horizontal/Vertical
- Stacked/side-by-side
- Standard/Logarithmic

Table 13-2: Visual Modifiers and Descriptions

Graph Type	Associated Icon	Purpose
Horizontal		Toggles the orientation of the data displayed between horizontal and vertical. This icon does not apply to pie chart data.
Stacked		Toggles the placement of data displayed between side-by-side and top-to-bottom. This icon does not apply to pie or line chart data.
Logarithmic scale		Toggles the numeric scale from standard increments to logarithmic. This icon does not apply to pie chart data.

Creating a Data Analysis Table

Create a Data Table

1. Click on one or more ports or interface cards from the Element Manager main screen. Hold down the <Ctrl> key to select multiple objects.
2. Select Report -> Graph Port... if the selected set of objects are ports. A graph window opens.
3. Select Report -> Graph Card if the selected object (s) are cards. A graph window opens.

4. If multiple ports (or cards) are selected, then select the type of counter data you want to display from the counter scroll-list on the bottom of the graph window.

The screenshot shows a window titled '10.3.102.66' with a tab labeled 'Interface'. It contains a table with the following data:

	Port 16/5	Port 16/9
InOctets	9,216,093,532	16,332
InUcastPkts	9,735,859	136
InMulticastPkts	0	0
InBroadcastPkts	0	0
InDiscards	0	0
InErrors	0	0
InUnknownProtos	0	0
OutOctets	8,852,105,132	36,840
OutUcastPkts	9,370,927	135
OutMulticastPkts	0	0
OutBroadcastPkts	0	0
OutDiscards	0	0
OutErrors	0	0

Below the table is a control bar with a scroll menu set to 'Absolute Value' and an 'Elapsed' timer showing '00:01:02'. An arrow points to the 'Absolute Value' dropdown menu with the label 'Select Data Field'.

5. Select the statistical data to be monitored.
 - a. Hold down the <Ctrl> key and click multiple fields to select specific counters
or
 - b. Left-click and drag the mouse to define a selection rectangle

Export a Data Table

Once you have created a data table, you can export the table to a saved file for future reference.

- Click the Export table icon.

The screenshot shows a window titled "10.3.102.66" with a tab labeled "Interface". Below the tab is a table with the following data:

	Port 16/4	Port 16/5
InOctets	20,569,193,844	9,287,097,116
InUcastPkts	21,961,166	9,821,721
InMulticastPkts	0	0
InBroadcastPkts	0	0
InDiscards	0	0
InErrors	0	0
InUnknownProtos	0	0
OutOctets	22,446,086,112	8,877,363,524
OutUcastPkts	26,838,055	9,412,035
OutMulticastPkts	0	0
OutBroadcastPkts	0	0
OutDiscards	0	0
OutErrors	0	0

Below the table is a toolbar with several icons. An arrow points to the icon representing a floppy disk, which is used for exporting the data table. To the right of the toolbar are buttons for "Close" and "Help", a refresh rate dropdown set to "10s", a sorting dropdown set to "AbsoluteValue", and a timer showing "Elapsed: 00:00:02".

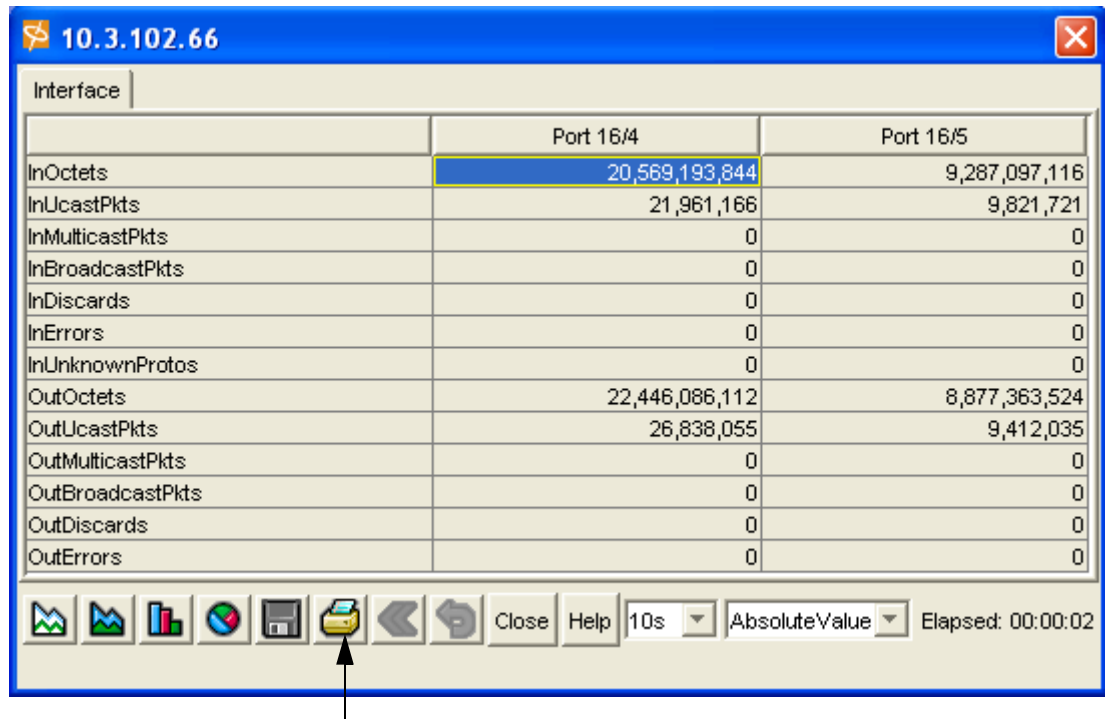
- Navigate to the location where you want to store the data table.
- Enter a file name for the data table.
- Save the file.

Print a Data Table

Once you have created a data table, you can print the table for reference.

- Create a data table.
Refer to [“Create a Data Table” on page 140](#).

- Click the print icon from the data table.



10.3.102.66

Interface

	Port 16/4	Port 16/5
InOctets	20,569,193,844	9,287,097,116
InUcastPkts	21,961,166	9,821,721
InMulticastPkts	0	0
InBroadcastPkts	0	0
InDiscards	0	0
InErrors	0	0
InUnknownProtos	0	0
OutOctets	22,446,086,112	8,877,363,524
OutUcastPkts	26,838,055	9,412,035
OutMulticastPkts	0	0
OutBroadcastPkts	0	0
OutDiscards	0	0
OutErrors	0	0

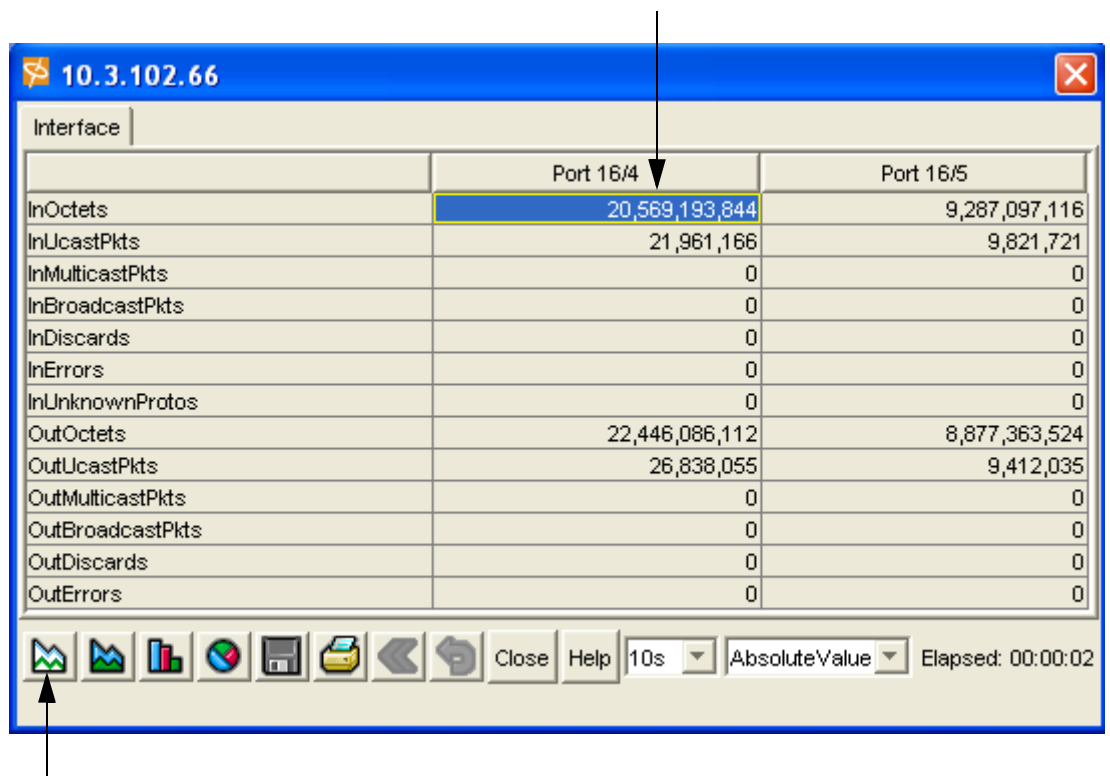
Toolbar: [Line Graph] [Area Graph] [Bar Graph] [Pie Chart] [Print] [Back] [Forward] [Close] [Help] [10s] [AbsoluteValue] [Elapsed: 00:00:02]

Creating a Data Analysis Graph


Statistical data is displayed in a table by default. However, you can use the data in the tabular format to create a variety of graphs.

- Create a table of network data.
Refer to [“Creating a Data Analysis Table” on page 140](#)

Once you have selected the data to be monitored, the Graph icons become active.



- Select the icon that represents the type of graph you want to create.
Refer to “Types of Graphs” on page 138
The graph appears in a new window.




Modify a Graph

All graphs (except for pie graphs) have the option of using visual modifiers to alter them. For example, a graph can be displayed horizontally or vertically.

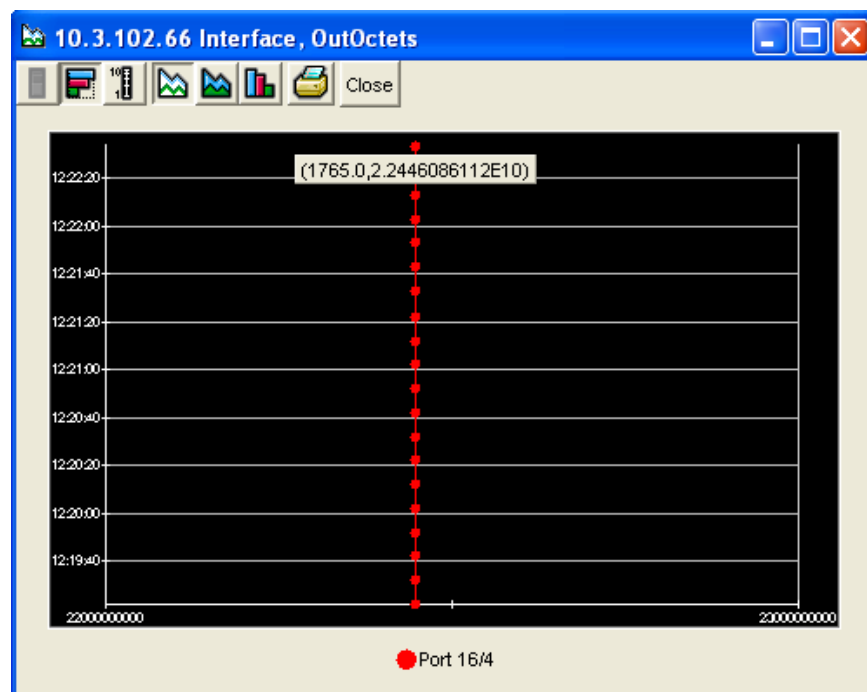
1. Click the icon that represents the visual modifier that you want to apply to the graph.

Refer to [“Visual Modifiers”](#) on page 140.

In the following example, the horizontal modifier was chosen.



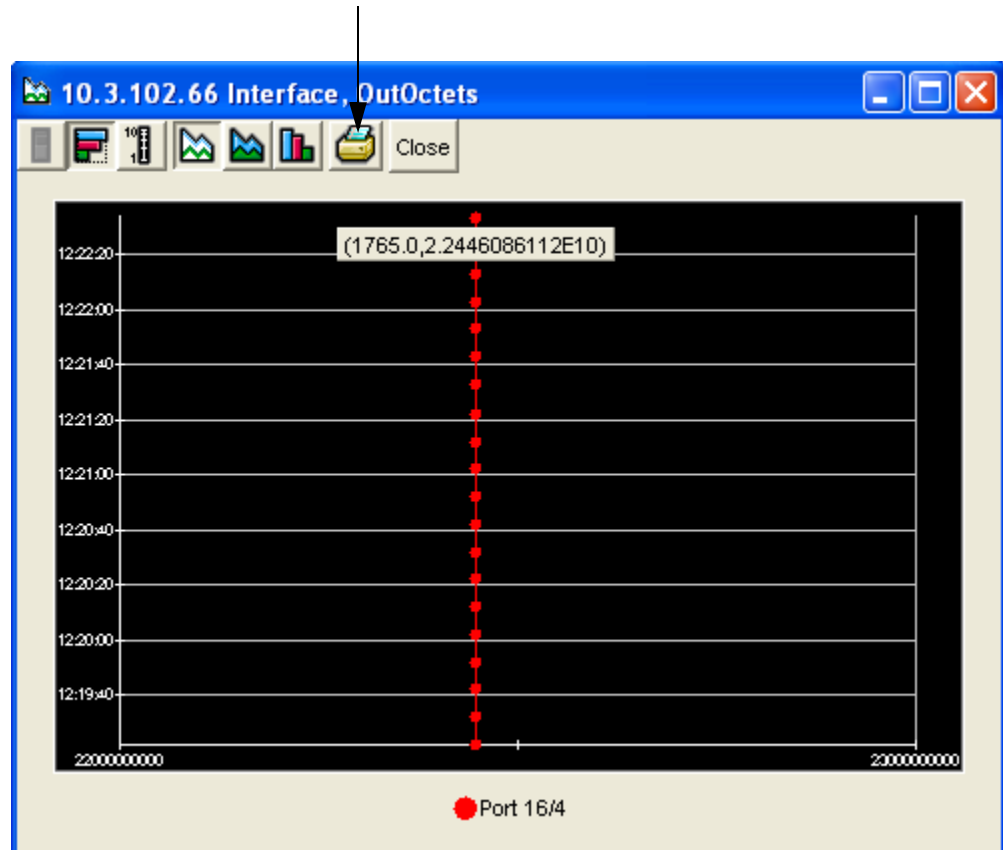
The graph is changed according the applied visual modifier.



Print a Graph

Once you have created a data graph, you can print the graph for reference.

1. Create a data table.
Refer to [“Create a Data Table” on page 140](#).
2. Click the print icon above the graph.



About SNMP Traps

Events Sent to Trap Receivers

Event messages are generated whenever a significant event occurs. Significant events sent to trap receivers include:

- Configuration changes
- Card insertion/removal
- Card up/down
- Port up/down
- Fan up/down
- Temperature problems
- InfiniBand Subnet Manager slave becomes a master (this indicates a reboot)
- InfiniBand Subnet Manager master becomes a slave (this indicates a reboot)
- InfiniBand Multicast group is added

- InfiniBand Multicast group is deleted
- InfiniBand Partition is added
- InfiniBand Partition is deleted
- InfiniBand Multicast member is added
- InfiniBand Multicast member is deleted
- InfiniBand Node is added to a partition
- InfiniBand Node is deleted from a partition
- InfiniBand Node is added to subnet
- InfiniBand Node is deleted from subnet

Configuring SNMP Settings

Viewing Current SNMP Trap Receivers

1. Launch Element Manager, if you have not already done so.
2. Select Health --> Trap Receivers.

The Trap Receivers window appears.



Any previously configured Trap Receivers appear in the window.

If the Receive Events field is set to false (as in the example above), the Receiver is not set to receive events.

Adding an SNMP Trap Receivers

1. Launch Element Manager, if you have not already done so.
2. Select Health --> Trap Receivers.

The Trap Receivers window appears.



3. Click the **Insert** button.

The Trap Receivers, Insert Trap... window appears.

4. Enter the IP address of the desired Trap Receiver in the **Address** field.
5. Enter the SNMP community string in the **Community** field.
6. Check the **Receive Events** box if you want to enable the Trap Receiver to receive events. You can also add the Trap Receiver, then enable it later.
7. Click the **Insert** button.

Editing a Current SNMP Trap Receiver

1. Launch Element Manager, if you have not already done so.
2. Select Health --> Trap Receivers.

The Trap Receivers window appears.

Address	Version	Community	Receive Events
172.16.1.34	v2c	public	false

1 row(s)

Any previously configured Trap Receivers appear in the window.

3. Click in any of the fields to edit them.
 - Click into the **Address** field to change the IP address of a Trap Receiver.
 - Click into the **Version** field to change the SNMP version. **Note:** v2C is the highest supported version.
 - Click into the **Community** field to change the SNMP community string.
 - Click into the **Receive Events** field to enable or disable the Trap Receiver from receiving events. The field becomes a drop-down menu when clicked.

Monitoring Through the CLI

This chapter gives an overview of the following:

- [“About InfiniBand Events” on page 149](#)
- [“About Tracing” on page 149](#)
- [“About SNMP Traps” on page 151](#)
- [“Configuring SNMP Settings” on page 152](#)

About InfiniBand Events

Event messages are generated whenever a significant event occurs in the system, and cannot be turned-off or configured by the user.

The following types of scenarios fall into the Events category.

- A Subnet Manager has been started (created)
- A new InfiniBand node is in-service
- An existing InfiniBand node is out-of-service
- A new InfiniBand multicast group is created
- A new multicast member is added to an existing multicast group
- An existing multicast group is deleted

About Tracing

Unlike Events, Trace messages are only generated when the user has explicitly enabling tracing for debugging/trouble-shooting purposes. By default, tracing is disabled.

The trace feature has different levels of control to allow you to dynamically select the amount and/or type of tracing information to be generated.

Types of Traces

The following types of traces are available:

- Application
- Module

Application

"Application" is used to specify the application for which tracing should be configured.

The amount of applications listed may change. The numbers that represent different applications may change between releases. Check application and module number assignments using CLI help before setting a trace level.

Module

"Module" is used to specify what module of code (library), in the selected application for which tracing should be configured.

The number of modules may change between releases. The numbers assigned to modules may also change. Check module number assignments using CLI help before setting trace levels.

Trace Levels

"Trace-Level" is used to specify the verbosity of the output. You can specify different levels of verbosity to control the amount of the tracing information that is generated.

When you select a level or tracing, the information at the specified level and below is shown.

For example, if you are only interested in seeing information that belongs to the VERBOSE category, you can dynamically set the control to VERBOSE. The trace software will then display information in the VERBOSE, TERSE and VERY_TERSE categories.

Trace levels are described in the following table:

Table 14-1: Trace Level Descriptions

Trace Level	Description
NO_DISPLAY	No tracing information will be generated. This is the default setting to ensure the best SM performance.
VERY_TERSE	Not currently used.
TERSE	This level is used to generate the basic flow of MAD packet as it goes through the Subnet Manager state machine. Verbose description will be made to the success or failure of the MAD packet. For example, a PathRecord lookup failure will include description on the cause of the failure (SGID is not found in database, DGID is not found in database, no component is specified in the component mask.) instead of simple success or failure code.
VERBOSE	This level shows the MAD management class header information.
VERY_VERBOSE	This level shows the MAD in a decoded format according to the decoding algorithm.
SCREAM	This level shows the MAD in its raw format.

Flow

"Flow" is used to specify a particular part of the code. The way that flows are specified varies depending on the module.

- 0x1 : This flow shows all the Subnet Management (SM) discovery, assignment routing and routing processing.
- 0x2: This flow shows all the Subnet Management Agent (SA) related processing.
- 0x4 : This flow shows the trap related processing.
- 0x8: This flow shows the redundancy related processing.
- 0x10: This flow shows the Related Multipack Protocol (RMPP) related processing.
- 0x20: This flow shows the partition management related processing.
- 0x1000: This flow shows the internal state machine (may create more flow for more granularity for path, mcast, ...).
- 0x2000: This flow shows the user configuration management related processing.
- 0x4000: This flow triggers the Subnet Manager to take a snap shot of the current logging information for the trace.

For example, use:

- 0x4001 to take a snapshot of the SM logging information.
- 0x4002 to take a snapshot of the SA information.
- 0x8000: This flow triggers the MAD packet to be displayed.

For example, use:

- 0x8001 to display SM MAD packets and SM related processing.
- 0x8002 to display MAD packets and SA related processing.

About SNMP Traps

Events Sent to Trap Receivers

Event messages are generated whenever a significant event occurs. Significant events sent to trap receivers include:

- Configuration changes
- Card insertion/removal
- Card up/down
- Port up/down
- Fan up/down
- Temperature problems
- InfiniBand Subnet Manager slave becomes a master (this indicates a reboot)
- InfiniBand Subnet Manager master becomes a slave (this indicates a reboot)
- InfiniBand Multicast group is added
- InfiniBand Multicast group is deleted
- InfiniBand Partition is added
- InfiniBand Partition is deleted
- InfiniBand Multicast member is added
- InfiniBand Multicast member is deleted

- InfiniBand Node is added to a partition
- InfiniBand Node is deleted from a partition
- InfiniBand Node is added to subnet
- InfiniBand Node is deleted from subnet

Configuring SNMP Settings

Viewing Current SNMP Trap Receivers

1. Enter the **show snmp** command in the CLI.
The current Trap Receivers are displayed in the output.

```

Topspin-360# show snmp
=====
                        SNMP Information
=====
                contact : Local HP support representative
                location : 3000 Hanover Street, Palo Alto, CA 94304
=====
                        Trap Receivers
=====
ipaddr          version      community      recv-events
-----
172.16.1.34    v2c             public         true

```

Add an SNMP Trap Receiver

1. Enter the **snmp-server host** command in global-configuration mode, as well as the following host information:
 - IP address or DNS name of an SNMP server.
 - SNMP community string that authenticates your device to the SNMP server.

```

Topspin-90> enable
Topspin-90# config
Topspin-90(config)# snmp-server host 10.3.106.99 secret

```

2. Store the contact information for your device by entering the **snmp-server contact** command, as well as the ASCII string of contact information.

```

Topspin-90(config)# snmp-server contact "Local HP support
Representative"

```

3. Store location information about your device by entering the **snmp-server location** command, as well as the ASCII text string of location information.

```

Topspin-90(config)# snmp-server location 3000 Hanover Street,
Palo Alto, CA 94304

```

Index

A

acceptable p_key values	56
access	
changing access levels	52
community strings	51
access levels	47
adding new users	49
admin	
user account commands	50
Authorized Trap Receivers	30
autoconnect, Element Manager	51
auto-negotiate	
set IB interface with CLI	40
set IB interface with GUI	34

B

base version	131
boot image	101
broadcasting messages to users	41

C

changing	
user access-levels	52
user identity	52
class version	131
CLI	
command completion	21
command history	22
ending a session	23
exiting command modes	20, 21
setting terminal parameters	22
understanding	20
clock	41
setting time (CLI)	42
cold sync timeout	69
cold synchronization	5
cold-sync-limit	86
cold-sync-period	86
cold-sync-timeout	86
command	
username	50
command completion	
using	21
example	22
command history	22

displaying	22
command modes	20
about	20
exiting	21
commands	
copy	99, 100
exit	21, 99, 100
gateway	99, 100
history	22
install	100
ip address	99, 100
no shutdown	99, 100
password	49
show boot-config	102
username	49, 50, 52
community strings	31
about	50
changing	51
element manager	51
configuration files	93, 105, 111
configuration, image, and log files	93, 105, 111
configure	
IB interface speed	
CLI	39
GUI	33
configuring	
Ethernet Management IP Address	10
system hostname	12
configuring SNMP	30
connecting hardware	95
Connectivity	
basic	10
establishing	10
contention	
SM routing	6
controller op cap	121
copy command	99, 100
creating	
hardware connection	95
new user account	49

D

database sync	67, 69
about	5
max-backup-sms	85
poll-interval	86
session-timeout	85
database synch	
master active sm	68
master poll interval	68

master poll retries	68
default link down	134
default partitions	54
default passwords	46
deleting	
a user account	50
image files	103
system images	103
destination LID	78
device ID	131
device ID field	131
device vendor ID	131
diag code	133
disable	
autoconnection with element manager	51
displaying	
configuration, image, and log files	106, 112
user information	48
distance	
SM routing	6
downloading	
a new system image	98
downloading image	99
E	
Element Manager	
starting	29
enabling	
no shutdown command	99, 100
ending a CLI session	23
Ethernet Management IP Address	
configuring	10
ethernet management port	95
events	
Trap Receivers	146, 151
exit command	21, 99, 100
F	
flow	
about	151
flow mask	151
flow label	74
FTP	98, 99
G	
gateway	
command	99, 100
general tab	73

GID prefix	133
global route header	74
global-configuration mode	20
GRH	74, 89
GUID	63, 83, 130, 132
view	130, 132
view the GUID	126
GUIDCap	135
GUIDs	
view	129
H	
hardware connection	95
Help	24
hex to binary conversions	55
history command	22
hop limit	74
HOQ life	135
hostname, customizing	12
hot standby	
about	5
I	
I/O controller units	120, 121
IB interface speed	
configure with CLI	39
configure with GUI	33
image	
download from TFTP server	99
image files	93, 105, 111
deleting	103
maximum number	103
understanding	93
in-band connection	95
InFilterRawPktEnf	135
infiniBand subnet details	
description field	130
GUID	130
node tab	130
show advanced button	130
SubnetPrefix field	130
type field	130
InPartEnforce	135
install	100
command	100
IOC	120, 121
IOC services	121
IOSubclass	120
ip address	99, 100

IPoIB	2	about	3
		supported implementations	3
L		MTU	74, 89, 135
LID	132, 133	multicast group	73
destination	78	members tab	74
source	78	N	
link speed active	134	neighbor MTU	135
link speed enabled	134	Network Connectivity	
link speed supported	134	testing	13
link state machine	133	new boot image	
link width enabled	134	specify	101
link width supported	134	new users	49
LinkSpeedSupported	134	new-session-delay	86
LMC	134	no shutdown	99, 100
local identifier	132	node GUID	130, 132
local identifier mask control	134	node information	
local port number	131	advanced	130
LocalOverrunError	136	basic	129
LocalPhyError	136	node type field	130
log files	93, 105, 111	nodes tab	130
about	111	numports field	130
logging		O	
about	149	oper VL	135
module	150	OutFilterRawPktEnf	135
login		out-of-band connection	95
telnet	11	OutPartEnforce	135
M		P	
M_Keys	135	p_key	
MAD	131	about	53
management datagram	131	acceptable values	56
master active sm		advanced port information	135
database sync	68	pkey violation	135
master poll interval		packet life time	74, 89
database sync	68	partition cap field	131
master poll retriesl		partitions	
database sync	68	acceptable p_key values	56
master sm sl	135	default	54
master subnet manager	81	selecting a p_key value	54
MasterSmLID	133	password	
max-backup-sms	85	about	46
members tab	74	default	46
MIBs	30	setting or changing	47
MKey	133	permission levels	47
MKey violation	135	physical ports	
MKeyLeasePeriod	134	NumPorts field	130
MLID	74, 89		
module	150		
MPI			

PKey	74, 89	SL	74, 90
poll interval		SM routing	
synchronization	69	contention	6
poll-interval	86	distance	6
port transition	133	SMP	135
PortGUID field	131	SNMP	
privileged-execute mode	20	adding new users	49
protocol version	120	changing community strings	51
		configuring	30
Q		element manager	51
Q_Key violation	135	source LID	78
Q_Keys	135	specifying new boot image	101
QKey	74	SSH	
queue pairs	3	login	11
		standby subnet manager	82
R		Starting the Element Manager	29
rate	74, 89	state	133
RDMA transfer size	121	storage manager	26
read/write access	47	sub-command mode	
rebooting		entering	21
about	37, 43	overview	21
system	102	subnet management agents	
recovery image	94	understanding	5
reload	43	subnet manager	
Remote SSH Login	15	activity count	63
Remote Telnet Login	14	basic node information	129
RespTime	135, 136	cold-sync-limit	86
resync-interval	86	cold-sync-period	86
roles and privileges	46	cold-sync-timeout	86
		details	129
S		discovering	63
SA	5	hot standby	5
scope	74	master	81
scream	150	max-backup-sms	85
selecting a p_key value	54	new-session-delay	86
send message que depth	120	notActive	63
Sending Messages to Individual Users	41	poll-interval	86
service ID	121	priority	63
service level	74, 90	response timeout	63
service name	121	resync-interval	86
services entries	121	session-timeout	85
session-timeout	85	smKey	63
setting		standby	82
time (CLI)	42	status	63
set-up procedures	9	sweep interval	63
show advanced button	130	subnet manager routing	
show boot-config		about	5
command	102	subnet timeout	135
		SubnetPrefix field	132
		subsystem ID	120
		subsystem vendor ID	120

super	46
switch element	
display route	91
switch elements	
sm routing	6
synchronization	
cold sync	5, 69
poll interval	69
transactional	5
syslog	114
system hostname	12

T

TClass	74, 89
telnet	
login	11
terminal parameters	
setting	22
tftp	
image upgrade summary	95
TFTP server	99
Topspin 360	
introducing	1
trace	
flow descriptions	151
module	150
trace level	
descriptions	150
no display	150
scream	150
terse	150
verbose	150
very terse	150
very verbose	150
transactional synchronization	5
transition	133
trap receivers	
events	146, 151
type of services	4

U

upgrading	
copy command	99, 100
reboot	102
specify boot image	101
user access	
changing	52
user access levels	
changing	52

user account	
configuration commands	50
user accounts	
adding	49
user information	
displaying	48
user-executive mode	20
usernames and passwords, understanding	46
Using	21

V

VL stall count	135
----------------------	-----

W

work queues	3
-------------------	---